

# SCIENTIFIC DATA



OPEN  
ANALYSIS

## An empirical meta-analysis of the life sciences linked open data on the web

Maulik R. Kamdar <sup>1,2</sup>✉ & Mark A. Musen <sup>1</sup>

While the biomedical community has published several “open data” sources in the last decade, most researchers still endure severe logistical and technical challenges to discover, query, and integrate heterogeneous data and knowledge from multiple sources. To tackle these challenges, the community has experimented with Semantic Web and linked data technologies to create the Life Sciences Linked Open Data (LSLOD) cloud. In this paper, we extract schemas from more than 80 biomedical linked open data sources into an LSLOD schema graph and conduct an empirical meta-analysis to evaluate the extent of semantic heterogeneity across the LSLOD cloud. We observe that several LSLOD sources exist as stand-alone data sources that are not inter-linked with other sources, use unpublished schemas with minimal reuse or mappings, and have elements that are not useful for data integration from a biomedical perspective. We envision that the LSLOD schema graph and the findings from this research will aid researchers who wish to query and integrate data and knowledge from multiple biomedical sources simultaneously on the Web.

### Introduction

Over the last decade, the biomedical research community has published and made available, on the Web, several sources consisting of biomedical data and knowledge: anonymized medical records<sup>1</sup>, imaging data<sup>2</sup>, sequencing data<sup>3</sup>, biomedical publications<sup>4</sup>, biological assays<sup>5</sup>, chemical compounds and their activities<sup>6</sup>, biological molecules and their characteristics<sup>7</sup>, knowledge encoded in biological pathways<sup>8,9</sup>, animal models<sup>10</sup>, drugs and their protein targets<sup>11</sup>, medical knowledge on organs, symptoms, diseases, and adverse reactions<sup>12</sup>. However, biomedical researchers still face severe logistical and technical challenges to integrate, analyze, and visualize heterogeneous data and knowledge from these diverse and often isolated sources. Researchers need to be aware of the availability and accessibility of these sources on the Web, and need extensive computational skills and knowledge to query and explore them. To systematically consume and integrate the data and knowledge from these sources for use in their analyses, researchers often spend a significant amount of their time dealing with the heterogeneous syntaxes, structures, formats, schemas, and biomedical entity notations used across these sources.

The complexity and time in inter-disciplinary scientific research drastically increase as the biomedical researcher ends up learning multiple systems, configurations, and means to reconcile similar entities across different sources. In most cases, the biomedical researcher just wishes to retrieve relevant data pertaining to their unique criteria or to retrieve answers to specific queries, such as “*What are the medications prescribed to melanoma patients who have a V600E mutation in their BRAF gene?*”, or “*List molecular characteristics of antineoplastic drugs that target Estrogen Receptor  $\alpha$  and have molecular weight less than 300 g/mol*”. However, due to the current state of the biomedical data and knowledge landscape with multiple, isolated, heterogeneous sources, the researcher has to hop across several Web portals (e.g., PubChem<sup>5</sup>) and search engines (e.g., PubMed<sup>4</sup>) for these tasks.

Semantic Web and linked data technologies are deemed to be promising solutions to tackle these challenges, and enable toward Web-scale data and knowledge integration and semantic processing in several biomedical domains, such as pharmacology and cancer research<sup>13–17</sup>. Biomedical researchers have been some of the earliest adopters of these technologies, and have used them to drive knowledge management, semantic search, clinical decision support, rule-based decision making, enrichment analysis, data annotation, and data integration<sup>12,16,18–20</sup>. Moreover, these technologies are used to represent data and knowledge stored in isolated, heterogeneous sources

<sup>1</sup>Center for Biomedical Informatics Research, Stanford University, Stanford, CA, USA. <sup>2</sup>Present address: Elsevier Health Markets, Philadelphia, PA, USA. ✉e-mail: [maulik@maulik-kamdar.com](mailto:maulik@maulik-kamdar.com)

on the Web to create a network of linked open data and knowledge graphs, often referred to as the Linked Open Data (LOD) cloud<sup>21</sup>. The LOD cloud was conceived from a vision that a decentralized, distributed, and heterogeneous data space, extending over the traditional Web (i.e., World Wide Web, or Web of Documents), can reveal hidden associations that were not directly observable<sup>22</sup>. The linked open data and knowledge sources, relevant to biomedicine, in the LOD cloud are collectively described as the Life Sciences Linked Open Data (LSLOD) cloud<sup>13</sup>. A diagrammatic representation of the LSLOD cloud can be found as part of the popular Linked Open Data cloud diagram<sup>23</sup> (<https://lod-cloud.net/>).

These technologies have matured enough that they are also being widely adopted by several companies, and subsets of the generated knowledge graphs are used in biomedical applications (e.g., Google<sup>24</sup>, IBM<sup>25</sup>, Elsevier<sup>26</sup>, NuMedii<sup>27</sup>, Pfizer<sup>28</sup>). Semantic Web technologies could also be combined with natural language processing-based relation extraction methods to systematically extract scientific relations between biomedical entities and to make them available for querying and consumption through a structured representation of the knowledge graph<sup>26,29</sup>.

Ideally, the LSLOD cloud should serve as an open Web-based scalable infrastructure through which biomedical researchers can query, retrieve, integrate, and analyze data and knowledge from multiple sources, without the requirement on the part of the researchers to download and manually integrate those sources and be concerned with the location, heterogeneous schemas, syntaxes, varying entity notations, representations, or the mappings to reconcile similar concepts, relations, and entities between these sources. However, it is well known that the LSLOD cloud has several shortcomings for this vision to be achieved. Specifically, most biomedical researchers find it difficult to query and integrate data and knowledge from the LSLOD cloud for use in their downstream applications, since the LSLOD cloud faces a set of challenges similar to those that it had set out to solve in the first place<sup>13,30–34</sup>. The rampant semantic heterogeneity across the sources in the LSLOD cloud, due to the heterogeneous schemas, the varying entity notations, the lack of mappings between similar entities, and the lack of reuse of common vocabularies, leads to the question on whether the LSLOD cloud can truly be considered “linked”.

In this paper, we present our results on an empirical meta-analysis conducted on more than 80 data and knowledge sources in the LSLOD cloud. The main contributions of this research can be outlined as follows:

1. Using an approach to automatically extract schemas and vocabularies from a set of publicly available LOD endpoints that expose biomedical data and knowledge sources on the Web, we have created an LSLOD schema graph encapsulating content from more than 80 sources.
2. We estimate the extent of reuse of content in LSLOD sources from popular biomedical ontologies and vocabularies, as well as the level of intra-linking and inter-linking across different LSLOD sources.
3. Combining methods of similarity computation and community detection based on word embeddings, we identify communities of similar content across the LSLOD cloud.

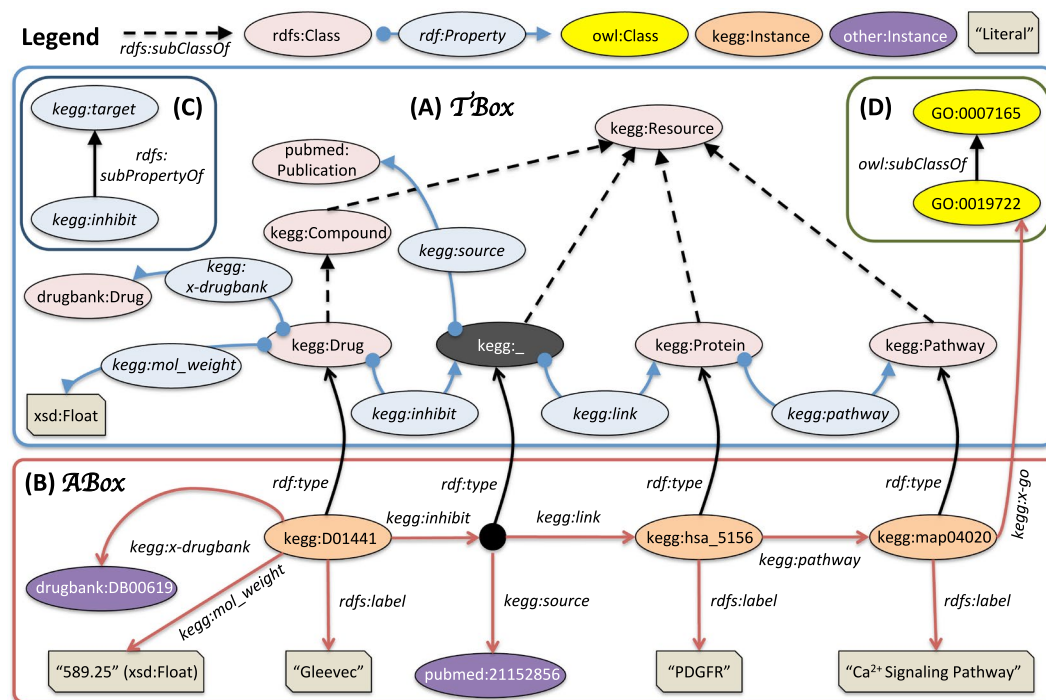
We believe that the resources and findings established through this meta-analysis will aid both: (i) the biomedical researchers, who wish to query and use data and knowledge from LSLOD cloud in their applications, and (ii) the Semantic Web researchers, who can develop methods and tools to increase reuse and reduce semantic heterogeneity across the LSLOD cloud, as well as develop efficient federated querying mechanisms over the LSLOD cloud to handle heterogeneous schemas and diverse entity notations. In subsequent sections of this paper, we provide a brief overview on the Semantic Web technologies and the LSLOD cloud, delve deeper into the different biomedical data and knowledge sources that were used in this research, and provide an overview on the different methods used to extract schemas and vocabularies from publicly available LOD graphs and to evaluate the reuse and similarity of content across the LSLOD cloud. We will present the results from this empirical meta-analysis of the quality, reuse, linking, and heterogeneity across the LSLOD cloud, summarize some of our key findings and observations on the state of biomedical LOD on the Web, and discuss what it means moving forward on the querying and consumption of data and knowledge on the Web in biomedical applications, in a more systematic and integrated fashion.

The LSLOD schema graph and other pertinent data, as well as the results and visualizations from this meta-analysis are made available online at <http://onto-apps.stanford.edu/lslodminer>. The code used in this research is made available on GitHub (<https://github.com/maulikkamdar/LSLODQuery>).

## Background and Related Work

**Semantic web technologies and the life sciences linked open data (LSLOD) Cloud.** To achieve the goals of a Linked Open Data (LOD) cloud over the Web, the Semantic Web community has developed several standards, languages, and technologies that aim to provide a common framework for data and knowledge representation — linking, sharing, and querying, across application, enterprise, and community boundaries. Semantic Web languages and technologies have been used to represent and link data and knowledge sources from several different fields such as life sciences, geography, economics, media, and statistics, to essentially create a linked network of these sources and to provide a scalable infrastructure for structured querying of multiple heterogeneous sources simultaneously, for Web-scale computation, and for seamless integration of big data.

Some Semantic Web languages and technologies have been standardized and recommended by the World Wide Web Consortium (W3C) and are also widely used by the biomedical community to create the LSLOD cloud. For example, The Resource Description Framework (RDF), a simple, standard triple-based model for data interchange on the Web, is used to represent information resources (e.g., biomedical entities, relations, classes) as linked graphs on the LSLOD cloud<sup>35</sup>. Each resource is uniquely represented using a HTTP Uniform Resource Identifier (URI), so that users can view information pertaining to that resource using a Web browser (i.e., Web-dereferenceable). An example of an HTTP URI from the LOD graph of DrugBank — a knowledge base of drugs and their protein targets<sup>11</sup> — is <http://bio2rdf.org/drugbank:DB00619>, where <http://bio2rdf.org/drugbank>: is considered to be the URI namespace — a declarative region that provides a scope to the identifiers, and



**Fig. 1** Diagrammatic representation to depict annotated RDF graphs: The KEGG RDF graph whose vocabulary is annotated using RDFS constructs and classes from the Gene Ontology. (a) The *T*-Box consists of the terminological component of the KEGG RDF Graph (i.e., the vocabulary). The pink nodes represent the different classes from KEGG and other RDF graph vocabularies (e.g., `drugbank:Drug`) annotated with `rdfs:Class` construct. The blue nodes represent the different properties annotated with `rdf:Property` (with associated domain and range classes). (b) The *A*-Box consists of the assertional component of the KEGG RDF graphs, with different individuals (e.g., `kegg:D01441`) explicitly typed with corresponding classes from *T*-Box. (c) The *R*-Box (relational component) consists of the property hierarchy. (d) The KEGG individual (pathway `Ca2+ SIGNALING PATHWAY`) may be cross-linked to the corresponding Gene Ontology class, causing a semantic mismatch (i.e., a class is considered equivalent to an individual).

DB00619 is the specific identifier of the drug Gleevec. RDF extends the linking structure of the Web by using the URIs to represent relations between two resources. Hence, RDF facilitates integration and discovery of relevant data and knowledge on the Web even if the schemas and syntaxes of the underlying data sources differ. The vocabularies and schemas of biomedical RDF graphs are generally annotated by the elements from the Resource Description Framework Schema (RDFS) language<sup>36</sup>. Similarly, in the last few years, the biomedical community has adopted the Web Ontology Language (OWL) as the consensus knowledge representation language to develop biomedical ontologies<sup>37</sup>. RDFS vocabularies and OWL ontologies enable the modeling of any domain in the world by explicitly describing existing entities, their attributes and the relationships among these entities. These entities and relationships are usually classified in specialization or generalization hierarchies. An ontology also contains logical definitions of its terms, along with additional relations, to facilitate advanced search functionalities<sup>38</sup>. Whereas, RDF is essentially only a triple-based, schema-less modeling language, semantics expressed in RDFS vocabularies and OWL ontologies can be exploited by computer programs, called reasoners, to verify the consistency of assertions in an RDF graph and to generate novel inferences<sup>13</sup>.

After the data and knowledge sources have been published as LOD using the RDF triple-based model, with semantics encoded using the RDFS and OWL languages, the SPARQL Protocol and RDF Query Language (SPARQL) can be used to query across multiple diverse sources using federated architectures<sup>17,39,40</sup>. While this distributed querying approach is inspired from the relational database community, SPARQL query federation architectures leverage the advantages provided by the graphical, uniform syntax, and schema-less nature of RDF to achieve query federation with minimal effort.

**Representative example on use of semantic web technologies in biomedicine.** Using Fig. 1, we give a more concrete real-world example on the use of Semantic Web technologies for representing data in the Kyoto Encyclopedia of Genes and Genomes (KEGG) — an integrated data source consisting of several databases, broadly categorized into biological pathways, proteins, and drugs<sup>9</sup> — and linking the different biomedical entities to similar entities in other biomedical sources, such as DrugBank<sup>11</sup>, PubMed<sup>4</sup>, and Gene Ontology — a unified representation of gene and gene product attributes across all species<sup>19</sup>.

`kegg:Drug`, `kegg:Protein`, and `kegg:Compound` are all RDFS classes in the KEGG vocabulary (Fig. 1a). It should be noted that `kegg:Drug` is a concise representation for a URI <http://bio2rdf.org/kegg:Drug>. `kegg:Drug` is a subclass of `kegg:Compound` (shown using `rdfs:subClassOf`). `kegg:inhibit` and

`kegg:target` are both RDFS properties, with `kegg:inhibit` being a sub-property of `kegg:target` (shown using `rdfs:subPropertyOf` in Fig. 1c). Using a reasoning-enabled query engine, a user who queries for `kegg:Compound` entities should retrieve all `kegg:Drug` entities, as well as all non-drug compounds. Similarly, all queries that look for `kegg:target` relations should retrieve all `kegg:inhibit` relations, along with non-inhibiting interactions. The domain and range of the properties are also shown using the  $\bullet$  and  $\blacktriangleright$  arrow heads respectively. The KEGG RDF Graph can be aligned to this KEGG RDFS vocabulary, with a class (e.g., `kegg:Drug`) instantiated with corresponding instances (e.g., `kegg:D01441` (GLEEVEC)) using the `rdf:type` construct, and properties realized as property assertions between different instances (e.g., `kegg:pathway` assertion between `kegg:hsa_5156` (PDGFR) and `kegg:map04020` (Ca<sup>2+</sup> SIGNALING PATHWAY)). These property assertions are shown using red arrows (Fig. 1b).

Figure 1b also shows an example of RDF *reification*, where a *blank node* (represented using a black colored node) is used to capture and represent additional information on the interaction between the drug GLEEVEC and the receptor protein PDGFR, involved in regulating cell growth and implicated in cancer. The *blank node* links the additional information to a publication with a URI from a PubMed LOD graph. Object properties (e.g., `kegg:pathway`) generally have `rdfs:Class` classes as domains and ranges, whereas data properties (e.g., `kegg:mol_weight`) link an `rdfs:Class` to a *datatype* (e.g., `xsd:Float` is the datatype for float numbers).

RDF graphs, annotated using the RDFS language, can be divided into two main components: the terminological component *T*-Box and the assertional component *A*-Box (Fig. 1). The *T*-Box consists of only the axioms pertaining to the vocabulary or schema (e.g. `rdfs:Class`, `rdfs:subClassOf` axioms, along with other logical combination constructs). The *A*-Box consists of axioms pertaining to the individuals and their property assertions. RDF graphs (and their associated vocabularies) can be exposed through SPARQL endpoints on the LSLOD cloud for programmatic access.

In general, classes, properties, and instances in RDF graphs may be cross-linked with elements from external ontologies and vocabularies, as well as with entities from other LOD graphs, and making such linkages is often considered a best practice when modeling and publishing datasets on the LOD cloud<sup>41</sup>. For example, `kegg:D01441` (GLEEVEC) of the KEGG RDF graph is linked to a similar entity `drugbank:DB00619` of the DrugBank RDF graph, and `kegg:map04020` (Ca<sup>2+</sup> SIGNALING PATHWAY) is mapped to the similar class `GO:0019722` in the Gene Ontology (Fig. 1b,d).

**Semantic heterogeneity across linked open data graphs.** Entity reconciliation and integrated querying are major problems in biomedicine, as there is often no agreement on a unique representation of a given entity or a class of entities<sup>41</sup>. Many biomedical entities are referred to by multiple labels, and the same labels may be used to refer to different entities. For example, similar gene entities could be represented using Ensembl<sup>42</sup>, Entrez Gene<sup>43</sup>, and Hugo Gene Nomenclature Committee (HGNC)<sup>44</sup> identifiers simultaneously.

One of the key principles for publishing LOD is the use of HTTP Uniform Resource Identifiers (URIs) for entity reconciliation and integrated querying. Each entity (e.g., LEPIRUDIN), or a class (e.g., Drug), should be represented uniquely using one unique URI and other publishers should reuse that URI in their datasets. However, most data publishers create their own URIs by combining the entity IDs from the underlying data sources (e.g., GLEEVEC DrugBank ID DB00619) with a unique namespace (e.g. <http://bio2rdf.org/drugbank:>) to represent entities. To resolve the problem of entity reconciliation, efforts such as Bio2RDF<sup>20</sup> and Linking Open Drug Data<sup>16</sup> have released guidelines for using cross-reference `x-ref` attributes rather than using the same URI. Similar entities in different sources should be mapped to each other (e.g., `drugbank:DB00619`  $\longleftrightarrow$  `kegg:D01441` to map GLEEVEC entities in DrugBank and KEGG), or all similar entities should be mapped to a common terminology (e.g., `drugbank:BE0000048`  $\longleftrightarrow$  `hgnc:3535`  $\longleftrightarrow$  `kegg:HSA_2147`, where the different URIs for protein PROTHROMBIN are mapped to the Hugo Gene Nomenclature Committee (HGNC) identifier) using such `x-ref` attributes from popular RDFS vocabularies<sup>41</sup>.

Ideally, in the case of RDF datasets, the terminological *T*-Box elements (i.e., classes and properties used in the schemas) should be reused from existing OWL ontologies or RDFS vocabularies. Large repositories of existing RDFS vocabularies are developed, either curated manually (e.g., Linked Open Vocabularies<sup>45</sup> — a catalog consisting of standardized versions of general RDFS vocabularies that are often used to encode additional semantics for RDF graphs on the LOD cloud) or by exploring the LSLOD cloud through automated crawlers. These repositories can serve as sources for data publishers to reuse elements from existing vocabularies. Moreover, in biomedicine, the classes from existing biomedical ontologies in the BioPortal repository<sup>46</sup> can also be reused when publishing LOD. RDF instances in the assertional *A*-Box may also be linked to classes from OWL ontologies using equivalence properties, such as `x-ref` or `owl:sameAs` as shown in Fig. 1d, where `kegg:x-go` and `kegg:x-drugbank` are subproperties of `x-ref`. However, cases where RDF ‘instances’ are mapped to ‘classes’ in other ontologies or LOD graphs using equivalence properties are often designated as *semantic mismatch* cases. An example of this scenario is depicted in Fig. 1b,d. The pathway Ca<sup>2+</sup> SIGNALING PATHWAY is represented as a class in the Gene Ontology (`GO:0019722`), but it is represented as an individual (instance) of the `kegg:Pathway` class in the KEGG RDF Graph (`kegg:map04020`). However, these two entities are linked and designated equivalent through the `kegg:x-go` property.

In some cases, the same relation may be expressed in different RDF graphs using different semantics or graph patterns<sup>13</sup>. For example, the same attribute of molecular weight is expressed in two different LOD sources, DrugBank and KEGG, using `drugbank:molecular-weight` and `kegg:mol_weight` data properties respectively. Ideally, the molecular weight attribute assertions should be structured using a uniform property across different sources. This property can originate from any of the LOV vocabularies or biomedical ontologies (e.g. CHEMINF:000216 “average molecular weight descriptor” from the Chemical Information Ontology<sup>47</sup>). As

a similar example, drug–protein target interactions are represented using different object properties and graph patterns across DrugBank (e.g., GLEEVEC  $\xrightarrow{\text{drug-target}}$  PDGFR) and KEGG (e.g., GLEEVEC  $\xrightarrow{\bullet}$  PDGFR, where the  $\bullet$  represents a *blank node* to capture additional information such as the source) RDF graphs. A user may wish to retrieve and reconcile such relations and attributes from both DrugBank and KEGG, as well as multiple other LOD graphs, to gain a complete picture in biomedicine, as unique drug–protein target relations may exist across the sources depending on the nature of their creation, or there may be conflicting information<sup>13,40</sup>. In the current state of the LSLOD cloud, the user must be aware of the underlying semantics and the data model used in these graphs as well as formulate lengthy SPARQL queries, which makes the entire process of information retrieval from the LSLOD cloud non-trivial.

There are several other benefits for reuse or cross-linking between similar entities or classes in different LOD graphs or ontologies<sup>41,48</sup> — *i*) reduction in data and knowledge engineering costs since the publisher can reuse existing minted URIs or link to existing entity descriptions, *ii*) enabling semantic interoperability among different datasets and applications, and *iii*) querying multiple LOD graphs simultaneously using federated architectures. The lack of reuse or cross-linking, the use of different semantics or graph patterns for modeling data in a given LOD graph, as well as semantic mismatch caused due to linking, are different manifestations of *semantic heterogeneity* across the LSLOD cloud.

**Related research.** Previously, we had conducted a study to estimate *term reuse* (i.e., reuse of classes from existing ontologies) and *term overlap* (i.e., similar classes across ontologies which are not reused) across biomedical ontologies stored in the BioPortal repository<sup>48</sup>. We found that ontology developers seldom reuse content from existing biomedical ontologies. However, there is significant overlap across the BioPortal ontologies. The concepts of *term reuse* and *term overlap* also hold true in the context of linked open data (LOD) sources. However, in the biomedical domain, there are several critical differences between ontologies and LOD sources due to the manner through which these sources are created and formalized. Some of these differences include: *i*) the vocabularies and schema for biomedical LOD sources are significantly smaller compared to most biomedical ontologies, *ii*) LOD sources have a much higher number of instances compared to biomedical ontologies and these instances may be reused from external resources, *iii*) LOD vocabularies formalized using the RDFS language may have less rich annotations compared to OWL ontologies (e.g., lack of synonyms and alternate labels), and *iv*) LOD sources, formalized using RDF and RDFS, may use *reification* to capture information at a higher granularity. Hence, the methods to analyze semantic heterogeneity across biomedical ontologies cannot be translated “as is” to analyze semantic heterogeneity across biomedical LOD sources.

A recent study found that the coverage of biomedical ontologies and public vocabularies is not sufficient to structure the entirety of biomedical LOD<sup>49</sup>. Hence, linked data publishers end up creating their own custom vocabularies, and do not publish them in a standardized way for reuse in other projects. Recently, multiple studies have been conducted on the larger Linked Open Data cloud to investigate the availability and licensing of Semantic Web resources, as well as to provide a mechanistic definition of what constitutes a ‘link’ in open knowledge graphs available on the Web<sup>30,33,50</sup>. However, these studies do not particularly focus on the semantic heterogeneity and vocabulary reuse across the LSLOD cloud<sup>51</sup>.

We use an Extraction Algorithm to extract and merge the schemas and vocabularies used by LSLOD sources for conducting our meta-analysis. The theory behind the Extraction Algorithm has identical characteristics to linked data profiling algorithms<sup>52–55</sup>. However, the linked data profiling algorithms have, in many cases, only been implemented over some of the popular SPARQL endpoints, do not account for variable SPARQL versions, require the retrieval of all instances and assertions that is often not scalable for LSLOD sources, or may not extend behind minimal schema extraction. Similarly, RDF graph pattern mining algorithms (e.g., evolutionary algorithms<sup>56</sup>) often require exhaustive training samples of (*source,target*) pairs which are often not possible to obtain for the LSLOD cloud.

## Methods

**Extracting schemas and vocabularies from the LSLOD cloud.** For conducting our meta-analysis of the Life Sciences Linked Open Data (LSLOD) cloud, we have selected a set of LOD projects, SPARQL endpoints, and RDF graphs, by querying the metadata of the projects in the “Life Sciences” section of the popular Linked Open Data cloud diagram (April 2018 version)<sup>23,30</sup>, as well as through a literature review<sup>13</sup> of articles, documenting popular biomedical LOD projects, available on the PubMed search engine<sup>4</sup>. Furthermore, we established the following criteria for an LSLOD source to be included in the meta-analysis: *(i)* Each LSLOD source must have a functional SPARQL endpoint, *(ii)* For cases when an LSLOD source does not have a functional SPARQL endpoint, the source should be available as RDF data dumps that can be downloaded and stored in a local SPARQL repository, and *(iii)* Each LSLOD source must have at least 1,000 instances under any classification scheme that can be queried through the SPARQL endpoint.

The entire list of the different LOD projects and details about sample RDF Graphs evaluated in this research are shown in Table 1. These RDF graphs include popular biomedical data sources such as DrugBank<sup>11</sup>, Kyoto Encyclopedia of Genes and Genomes (KEGG)<sup>9</sup>, Pharmacogenomics Knowledgebase (PharmGKB)<sup>57</sup>, Comparative Toxicogenomics Database (CTD)<sup>58</sup>, exposed by the Bio2RDF consortium<sup>20</sup>, UniProt<sup>7</sup>, ChEMBL<sup>59</sup>, and Reactome<sup>8</sup>, exposed by the EBI-RDF consortium<sup>60</sup>, and other sources such as The Cancer Genome Atlas (TCGA)<sup>3</sup>, PubChem<sup>5</sup>, National Library of Medicine (NLM) Medical Subject Headings (MeSH), National Bioscience Database Center (NBDC), and WikiPathways<sup>61</sup>. It should be noted that some of these projects (e.g., Bio2RDF<sup>20</sup>) have reprocessed publicly available data and knowledge published by other groups (e.g., DrugBank<sup>11</sup>), whereas in some cases data publishers themselves use Semantic Web technologies to provide access to their proprietary data as LOD (e.g., the European Bioinformatics Institute (EBI) RDF Platform<sup>60</sup>). These RDF graphs

Linked Data Project	Example RDF Graphs and Descriptions
Bio2RDF <sup>20</sup>	Several reprocessed data and knowledge sources such as:
	<b>DrugBank</b> <sup>11</sup> — Drug and drug target information
	<b>KEGG</b> <sup>9</sup> — Biological pathways, proteins, and drugs
	<b>PharmGKB</b> <sup>57</sup> — Protein–drug–disease relations
EBI-RDF <sup>60</sup>	Several proprietary data and knowledge sources such as:
	<b>Reactome</b> <sup>8</sup> — Biochemical reactions and pathways
	<b>BioSamples</b> <sup>70</sup> — Metadata about biological samples
	<b>ChEMBL</b> <sup>59</sup> — Drug-like molecules and activities
UniProt <sup>7</sup>	UniProt database on proteins and their characteristics
MO-LD <sup>94</sup>	Model Organism databases on different species such as Fly, Mouse, Yeast, Rat, Human and Zebrafish
NBDC RDF <sup>95</sup>	Japanese Life Sciences databases on chemicals, structures, sequencing data, pathways, and diseases
Linked Life Data <sup>91</sup>	Several reprocessed data and knowledge sources such as:
	<b>EntrezGene</b> <sup>43</sup> — Genes and variants in diseases
	<b>BioGrid</b> <sup>96</sup> — Protein and chemical interactions
	<b>IntAct</b> <sup>97</sup> — Protein–protein interactions
Linked TCGA <sup>17</sup>	DNA methylation, gene expression and clinical data of cancer patients from The Cancer Genome Atlas <sup>3</sup>
NIH PubChem <sup>5</sup>	Databases on substances, compounds, structures, and biological assays stored in the PubChem repository
WikiPathways <sup>61</sup>	Database of biological pathways maintained using a crowd-sourcing architecture
PathwayCommons <sup>98</sup>	Data warehouse consisting of several biological pathways and molecular interactions datasets
NLM MeSH <sup>86</sup>	Medical Subject Headings used to index biomedical publications in MEDLINE repository <sup>67</sup>
DisGenet <sup>99</sup>	Data warehouse on genes and variants associated to diseases
NextProt <sup>100</sup>	Data warehouse on protein structures and interactions
LinkedSPL <sup>101</sup>	Structured product labels of drugs represented as RDF

**Table 1.** Linked Open Data sources queried to generate the  $\mathcal{G}_{LSLOD}$  schema graph of the LSLOD cloud.

contain data and information on several biomedical entities such as drugs, genes, proteins, pathways, diseases, protein–protein interactions, drug–protein interactions, biological assays, and high-throughput sequencing. The SPARQL endpoint locations for these projects are listed at <http://onto-apps.stanford.edu/lslodminer>.

The meta-analysis relies on a set of extractors that navigate to each SPARQL endpoint and extract the schemas and vocabularies used in the RDF graphs exposed through those SPARQL endpoints. The extractors encapsulate an Extraction Algorithm that uses a set of SPARQL query templates (Box 1) to extract schemas (e.g., classes, properties, domains, ranges) as well as sample instances of those classes and their property values from the SPARQL endpoints in the LSLOD cloud. A snippet of the extracted information for a given source is shown in Box 2. The input to the Extraction Algorithm is the SPARQL endpoint, whose version of SPARQL can be automatically detected based on the SPARQL queries and keywords supported. The Extraction algorithm not only accounts for different versions of SPARQL supported at the remote SPARQL endpoints, but also uses alternative SPARQL queries if the remote endpoint times out during the extraction phase. The Extraction algorithm is presented in Box 3.

An extractor selects the set of SPARQL query templates for execution against the input SPARQL endpoint, given its version (**Step 2**). The extractor queries the SPARQL endpoint and extracts the named RDF Graphs  $\mathcal{N}\mathcal{G}$  exposed through the SPARQL endpoint (**Step 3, SPARQL Query SQ<sub>1</sub>**). For each RDF graph, the extractor extracts the set of classes  $\mathcal{C}$  and the total number of instances for each class  $k(\mathcal{C})$  (**Step 7, SQ<sub>2</sub>**). The extractor identifies the properties  $\mathcal{P}$  for which a specific class serves as the domain for all assertions, and also retrieves the range classes  $r(\mathcal{P})$  for object properties, or the datatypes  $dt(\mathcal{P})$  for data properties (**Step 12, SQ<sub>3</sub>**). The Extraction algorithm also extracts a random sample (without replacement) of 2,000 instances for each class (**Step 9, SQ<sub>4</sub>**) and 2,000 property assertions (**Step 14, SQ<sub>5</sub>**) for a class–property combination, and generates summary characteristics (e.g. is\_categorical, datatype, namespace, median length) of the associated data.

Classes, properties, and datatypes, are represented as nodes in the schema graph  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  extracted from the SPARQL endpoint by the specific extractor (**Steps 11, 17, 18, 21, 22**). The edges link property nodes to class nodes and datatype nodes depending on the domain and range descriptions of the corresponding property (**Steps 19 and 23**). The nodes and edges are annotated by the count ( $k(\mathcal{C})$  and  $k(\mathcal{P})$ ) and summary characteristics of corresponding instances or property assertions respectively (**Steps 11, 19, 23**). This schema graph is a fragment of the actual  $\mathcal{T}$ -Box of the RDF graph. The  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  graphs extracted for all RDF graphs in the LSLOD cloud are merged to generate the schema graph  $\mathcal{G}_{LSLOD}$  of the entire LSLOD cloud. A diagrammatic representation of this method and a subset of the merged  $\mathcal{G}_{LSLOD}$  schema graph are shown in Fig. 2.

**Determining vocabulary reuse across LSLOD sources.** Estimating the extent of observed vocabulary reuse across different LSLOD sources requires identifying the classes and properties most commonly reused from standard ontologies and vocabularies. For this goal, along with the  $\mathcal{G}_{LSLOD}$  schema graph extracted from the LSLOD cloud, we use two different corpora of ontologies and vocabularies:

**Box 1 SPARQL Query templates to extract schema descriptions from the LSLOD Cloud.**

```

----- SPARQL QUERY 1 -----
SELECT DISTINCT ?g WHERE {
    GRAPH ?g { ?s ?p ?o }
}

----- SPARQL QUERY 2 -----
SELECT ?Concept (COUNT (?x) AS ?cCount) WHERE {
    GRAPH <GRAPH_URI> { ?x rdf:type ?Concept }
} GROUP BY ?Concept ORDER BY DESC (?cCount)

----- SPARQL QUERY 3 -----
SELECT DISTINCT ?p ?c (COUNT (?x) AS ?count) ?valType WHERE {
    GRAPH <GRAPH_URI> { ``?x rdf:type <CONCEPT_URI>; ?p ?o.
    OPTIONAL { ?o rdf:type ?c } .
    FILTER (! (?p = ' rdf:type ')) .
    BIND (DATATYPE (?o) AS ?valType) }
} GROUP BY ?p ?c ?valType ORDER BY DESC (?count)

----- SPARQL QUERY 4 -----
SELECT ?x WHERE {
    GRAPH <GRAPH_URI> { ?x rdf:type <CONCEPT_URI> }
} ORDER BY RAND () LIMIT 2000

----- SPARQL QUERY 5 -----
SELECT ?x WHERE {
    GRAPH <GRAPH_URI> { ?c rdf:type <CONCEPT_URI>; <PROPERTY_URI> ?x }
} ORDER BY RAND () LIMIT 2000

```

**Box 2 Information from the Bio2RDF DrugBank RDF Graph extracted by the Extraction Algorithm.**

```

Source: DrugBank
Source URI: http://bio2rdf.org/drugbank\_resource:bio2rdf.dataset.drugbank.R3
Classes:
  Class: Drug
    Class URI: http://bio2rdf.org/drugbank\_vocabulary:Drug
    Sample Instances: ['http://bio2rdf.org/drugbank:DB03536', ...]
  Class: Drug-Drug-Interaction
    Class URI: http://bio2rdf.org/drugbank\_vocabulary:Drug-Drug-Interaction
  ...
Properties:
  Object Property: target
    Property URI: http://bio2rdf.org/drugbank\_vocabulary:target
    Property Realization: drugbank:Drug -> drugbank:target
      Domain: Drug, Range: Target
      Sample Assertion Values: ['http://bio2rdf.org/drugbank:BE0000059', ...]
  Data Property: molecular-weight
    Property URI: http://bio2rdf.org/drugbank\_vocabulary:molecular-weight
    Property Realization: drugbank:Enzyme -> drugbank:molecular-weight
      Domain: Enzyme, Range: Float
      Sample Assertion Values: [52221.099999999999, 56345.0, 57530.0, ...]

```

- BioPortal repository of biomedical ontologies:** BioPortal repository is the world's largest open repository of biomedical ontologies<sup>46</sup>. We obtained a dump of 685 distinct biomedical ontologies, serialized using the N-triples format and versioned up to January 1, 2018. This dump did not contain some ontologies that were deprecated or merged with existing ontologies, or added to BioPortal after January 1, 2018. This corpus includes ontologies such as Gene Ontology (GO)<sup>19</sup>, National Cancer Institute Thesaurus (NCIT)<sup>15</sup>, Chemical Entities of Biological Interest (ChEBI) Ontology<sup>6</sup>, and Systematized Nomenclature of Medicine - Clinical Terms (SNOMED CT)<sup>62</sup>, which are widely used in biomedicine for data annotation, knowledge management, and clinical decision support.
- Linked Open Vocabularies (LOV):** LOV is a catalog of RDFS vocabularies available for reuse with the aim of describing data on the Web<sup>45</sup>. Popular examples of RDFS vocabularies include the W3C Provenance

**Box 3 LSLOD Schema Extraction Algorithm.**

```

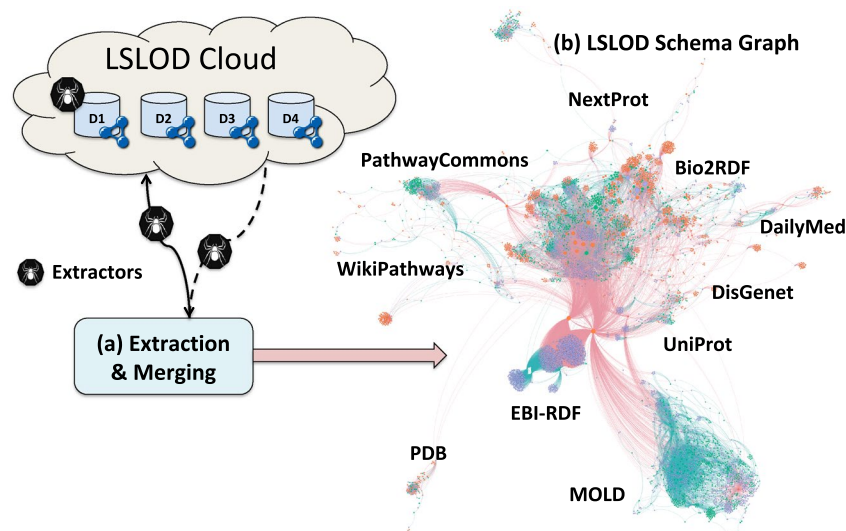
function EXTRACTOR ( $\mathcal{SE}$ ,  $v(\mathcal{SE})$ );
Input : A SPARQL Endpoint URI  $\mathcal{SE}$  and the supported version  $v(\mathcal{SE})$ 
Output: A dictionary of (graph URIs, schema ( $\mathcal{T}$ -Box) fragments, each represented as a
graph  $\mathcal{G}_{\mathcal{T}\text{-Box}}$ ), extracted from the SPARQL endpoint  $\mathcal{SE}$ 
/* Generate SPARQL Queries given the version of the endpoint (Listing 1) */
1  $\{SQ_i | i = 1, 2, \dots, 5\} \leftarrow \text{SPARQL\_QUERIES}(v(\mathcal{SE}))$ ;
/* Extract all the Graph URIs */
2  $\{URI_g\} \leftarrow SQ_1(\mathcal{SE})$ ;
/* An empty dictionary to store the Graph URIs and T-Box schema fragments */
3  $SF \leftarrow \{\}$ ;
4 for  $uri_g \in \{URI_g\}$  do
5    $\mathcal{G}_{\mathcal{T}\text{-Box}} \leftarrow \emptyset$ ;
6    $\mathcal{C}, k(\mathcal{C}) \leftarrow SQ_2(\mathcal{SE}, uri_g)$ ;
7   for  $c \in \mathcal{C}$  do
8     /* Extract a randomized sample of 2000 instances */
9      $\mathcal{A}_{(c)-2000} \leftarrow SQ_4(\mathcal{SE}, uri_g, c)$ ;
10    /* Generate a data summary of the instances */
11     $sum_c \leftarrow \text{SUMMARIZE}(\mathcal{A}_{(c)-2000})$ ;
12    Add node to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  with URI  $c$ , Type class, Count  $k(c)$  and Summary  $sum_c$ ;
13     $\mathcal{P}, r(\mathcal{P}), dt(\mathcal{P}), k(\mathcal{P}) \leftarrow SQ_3(\mathcal{SE}, uri_g, c)$ ;
14    for  $p \in \mathcal{P}$  do
15      /* Extract a randomized sample of 2000 assertions */
16       $\mathcal{A}_{(c,p)-2000} \leftarrow SQ_5(\mathcal{SE}, uri_g, c, p)$ ;
17      /* Generate a data summary of the assertions */
18       $sum_{cp} \leftarrow \text{SUMMARIZE}(\mathcal{A}_{(c,p)-2000})$ ;
19      /* Data property if range URI is missing, else Object property */
20      if  $r(p) = \emptyset$  then
21        Add node to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  with URI  $p$ , Type data property;
22        Add node to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  with URI  $dt(p)$ , Type datatype;
23        Add edges  $c \rightarrow p$  and  $p \rightarrow dt(p)$  to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$ , with Count  $k(p)$  and Summary
24         $sum_{cp}$ ;
25      else
26        Add node to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  with URI  $p$ , Type object property;
27        Add node to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$  with URI  $r(p)$ , Type class;
28        Add edges  $c \rightarrow p$  and  $p \rightarrow r(p)$  to  $\mathcal{G}_{\mathcal{T}\text{-Box}}$ , with Count  $k(p)$  and Summary
29         $sum_{cp}$ ;
30      end
31    end
32  end
33   $SF[uri_g] \leftarrow \mathcal{G}_{\mathcal{T}\text{-Box}}$ ;
34 end
35 return  $SF$ ;

```

Interchange<sup>63</sup>, Simple Knowledge Organization System (SKOS)<sup>64</sup>, and Schema.org<sup>65</sup>. We used 647 vocabularies that were present in the LOV catalog, as of January 1, 2018. The vocabularies in this catalog generally include concept descriptions (i.e., human-readable labels using the `rdfs:label` annotation properties).

For each URI in the  $\mathcal{G}_{\text{LSLOD}}$  schema graph (class, property, or sample instance URI), the *origin* of the URI is determined using a heuristic approach<sup>66</sup> by converting each URI to lowercase and using *regular expression* filters constructed from namespaces (e.g., `ncicb.nci.nih.gov` for NCIT<sup>15</sup>) and common identifier patterns.





**Fig. 2** Diagrammatic representation of the Extraction Algorithm: During the extraction phase, automated extractors (represented as black spiders here) query each SPARQL endpoint in the LSLOD cloud and extract the schemas (i.e., classes, object properties, data properties, URI patterns, graph patterns) to generate the LSLOD Schema Graph  $\mathcal{G}_{LSLOD}$ . A portion of  $\mathcal{G}_{LSLOD}$  is visualized here, using a force-directed network, with classes (purple nodes), object properties (green nodes) and data properties (orange nodes) extracted from multiple sources, such as Bio2RDF, EBI-RDF, etc.

For each LSLOD source, we compute the following statistics:

1. The number of ontologies and vocabularies reused in the source
2. The percentage of schema elements reused from external LSLOD sources
3. The percentage and distribution of classes whose entities are mapped or linked to other entities in external LSLOD sources (inter-linking)
4. The percentage and distribution of classes whose entities are mapped or linked to entities of other classes in the same LSLOD source (intra-linking)
5. The percentage of entities reused or mapped to external LSLOD sources

Previously, we outlined a method to estimate reuse of classes across multiple biomedical ontologies<sup>48</sup>. We repurposed this method to estimate vocabulary reuse across biomedical LOD sources. We generated an undirected network composed of different URIs from the  $\mathcal{G}_{LSLOD}$  schema graph as nodes (i.e., the classes and the properties), and the edges indicating reuse and mapping between the different schema element URIs. “Patterns” of instance URIs are also represented as nodes in this network to check if instances are reused or mapped to classes in biomedical ontologies (i.e., semantic mismatch). We extract connected components from this network and use Eq. 1 to estimate vocabulary reuse.

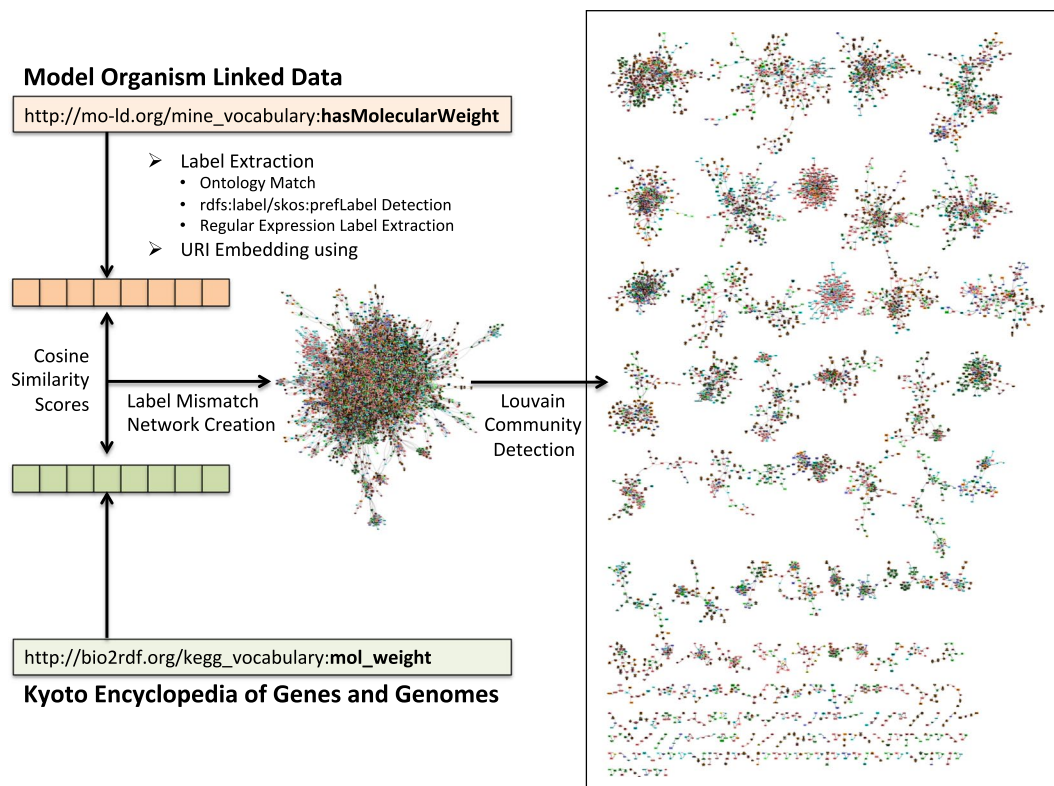
$$\text{Vocabulary Reuse} = \frac{\sum_{j|T_j \in \mathcal{M}_t} n_j - k}{N} \quad (1)$$

In the above equation, assume that the network is composed of  $k$  connected components  $\{T_0, T_1, \dots, T_k\}$ , and each component  $T_j$  is formed from  $n_j$  schema elements (i.e.,  $\{t_{0j}, \dots, t_{nj}\} \in T_j$ ). Assume that  $N$  is the total number of schema element URIs in the network. The number of terms in a component  $T_j$  must follow  $1 < n_j < N$  (i.e., components with a single term are not allowed). All schema elements in one component exhibit some reuse with respect to each other.

**Detecting communities of related content in the LSLOD cloud.** To detect clusters of related content in the LSLOD cloud, we used a 5-step algorithm that maps similar schema elements (e.g., classes and properties) between two different sources in the LSLOD cloud. A diagrammatic representation of this algorithm is shown in Fig. 3. This algorithm uses the schema elements extracted from the various LSLOD public SPARQL endpoints (see 1) and included in the  $\mathcal{G}_{LSLOD}$  graph. These schema elements may or may not be reused from a common vocabulary or ontology (vocabulary reuse in LSLOD cloud was determined using the methods described in the previous section).

The algorithm uses two sources to generate the mappings between two different schema elements:

1. **The set of biomedical ontologies and vocabularies that are reused in the LSLOD cloud:** Biomedical ontologies and vocabularies whose elements are reused across multiple RDF graphs were determined using the previous methods.



**Fig. 3** Diagrammatic representation of the algorithm to detect communities of similar content across LSLOD cloud: The labels for two URIs are extracted either (i) by matching to a reused ontology, or (ii) by using annotation properties such as *rdfs:label* or *skos:prefLabel*, or (iii) by using regular expressions on the URI. URI embeddings are generated from the extracted labels, by using a background source of MEDLINE biomedical abstracts. A similarity network is created with nodes representing the different URIs from the  $\mathcal{G}_{LSLOD}$  schema graph and edges representing the cosine similarity scores between two URI embeddings. We detect communities in this network using the Louvain method. The trivial example shown here is related to similar “molecular weight” URIs between the MO-LD project and the Bio2RDF KEGG LOD graph.

2. **Pre-trained word embedding vectors:** We use 100-dimensional word embedding vectors pre-trained from a corpus of approximately 30 million biomedical abstracts, stored in the MEDLINE repository<sup>67</sup>, using the GloVe (Global Vectors for Word Representation) method<sup>68</sup>. Word embeddings represent words as high-dimensional numerical vectors based on co-occurrence counts of those words as observed in the abstracts. The pre-trained embedding vectors and inverse document frequency (IDF) statistics for a vocabulary of 2,531,989 words are available under a CC BY 4.0 license at figshare<sup>69</sup>. These word embedding vectors have been used to map metadata fields in the BioSamples repository<sup>70</sup> to ontology terms in the BioPortal repository<sup>71</sup>, as well as to identify biomedical relations in unstructured text<sup>26</sup>.

**Label extraction.** In the first sweep, the algorithm extracts a label for each schema element URI (i.e., class, object property or data property) in the  $\mathcal{G}_{LSLOD}$  graph. This is done through three approaches in a sequential flow: [noitemsep]

1. A label is extracted from the values for the *rdfs:label*<sup>36</sup> or *skos:prefLabel*<sup>64</sup> annotation properties, if the  $\mathcal{G}_{LSLOD}$  schema element has already been reused from existing biomedical ontologies and vocabularies (see **Datasets**).
2. A label is extracted by querying the source (i.e., the SPARQL endpoint) of the schema element for *rdfs:label* or *skos:prefLabel* annotation assertions.
3. A label is generated using Regular Expression (*RegExp*) filters from the URI of the schema element. Specifically, the filters deal with cases where separators (e.g., “-”, “\_”) and camel case (e.g. *hasMolecularWeight*) is present in the URI. In Fig. 3, the labels “Has Molecular Weight” and “Mol Weight” will be extracted from the two URIs respectively.

**URI embedding generation.** We use the pre-trained word embedding vectors to represent each schema element URI in a high dimensional space. Specifically, a URI embedding vector is generated by computing a weighted average of the embedding vectors of the words in the URI label, with the weights being the IDF statistic for each word. A default embedding vector and IDF statistic is created for any word that is not present in the

vocabulary. The equation to generate a URI embedding vector is shown below. Here,  $\mathbf{x}_{(w_i)}$  represents the 100-dimensional word embedding vector, and  $\mathcal{L}(URI)$  is the URI label.

$$\mathbf{x}_{(URI)} = \frac{\sum_{w_i \in \mathcal{L}(URI)} \text{idf}_{(w_i)} * \mathbf{x}_{(w_i)}}{\sum_{w_i \in \mathcal{L}(URI)} \text{idf}_{(w_i)}} \quad (2)$$

**Cosine similarity score computation.** For any two schema element URIs that are not present in the same source, cosine similarity scores are computed using their URI embedding vectors. The equation to compute the similarity between two elements is presented below.

$$\text{Sim}(A, B) = \frac{\mathbf{x}_{(URI_A)} \cdot \mathbf{x}_{(URI_B)}}{\|\mathbf{x}_{(URI_A)}\| \|\mathbf{x}_{(URI_B)}\|} = \frac{\sum_{i=1}^{n=100} x_{(URI_A)i} * x_{(URI_B)i}}{\sqrt{\sum_{i=1}^{n=100} x_{(URI_A)i}^2} \sqrt{\sum_{i=1}^{n=100} x_{(URI_B)i}^2}} \quad (3)$$

**Creating a similarity network.** A similarity network is created where the different schema element URIs represent the nodes of this network. An edge exists between two nodes if: (i) the cosine similarity score between the corresponding URIs is greater than the threshold of 0.75, and (ii) there exists a one-to-one mapping between the two schema elements for the particular combination of LOD sources.

For example, consider three URI nodes  $n_1, n_2, n_3$ , such that,  $n_1, n_2 \in \mathcal{LD}_1$  and  $n_3 \in \mathcal{LD}_2$ . If  $\text{Sim}(n_1, n_2) = 0.87$ ,  $\text{Sim}(n_1, n_3) = 0.93$  then an edge will only be created between  $n_2$  and  $n_3$  nodes. This criteria has an underlying assumption that there should exist a unique alignment between similar schema elements in different LOD sources. While the cosine similarity scores in the previous step can, by itself, aid in detection of similar schema elements (e.g., all schema elements pertaining to *Molecular weight*), this similarity network will aid in the detection of sources that contain relevant information pertaining to a given biomedical concept or relation type.

**Community detection using Louvain method.** Finally, the Louvain method<sup>72</sup> for community detection in weighted networks is used. The cosine similarity scores form the edge weights in the similarity network. This method has been used to detect communities of users using social networks and mobile phone usage data, as well as detect species in network-based dynamic models<sup>73-75</sup>. The Louvain method maximizes a ‘modularity’ statistic  $Q$ , which measures the density of edges inside communities compared to edges between communities. In the equation below (reproduced from Blondel, *et al.*<sup>72</sup>),  $2m$  is the sum of all edge weights in the graph (i.e.,  $m = \frac{1}{2} \sum_{ij} \text{Sim}(i, j)$ ) and  $k_i$  is the sum of all weights to edges attached to node  $i$  (i.e.  $k_i = \sum_j \text{Sim}(i, j)$ ).  $\delta(c_i, c_j)$  is a delta function that looks at the community assignments ( $c_i$  and  $c_j$ ) for connected nodes  $i$  and  $j$ .

$$Q = \frac{1}{2m} \sum_{ij} \left[ \text{Sim}(i, j) - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j) \quad (4)$$

$$\delta(c_i, c_j) = \begin{cases} 1 & \text{if } (c_i = c_j) \\ 0 & \text{if } (c_i \neq c_j) \end{cases} \quad (5)$$

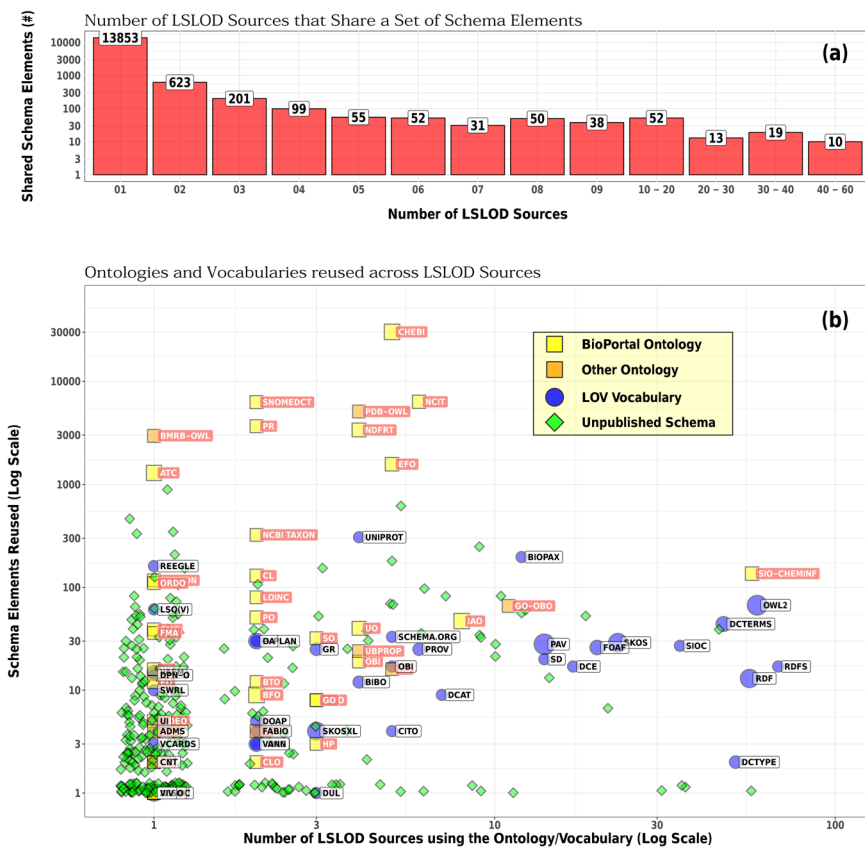
The communities are visualized using the Cytoscape platform<sup>76</sup>.

## Results

**LSLOD schema graph  $\mathcal{G}_{LSLOD}$ .** Using the Extraction Algorithm, we extracted 99 RDF graphs from  $\approx 20$  SPARQL endpoints. The total number of classes, object properties, data properties, and datatypes extracted and linked in the  $\mathcal{G}_{LSLOD}$  schema graph were 57,748, 4,397, 8,447 and 24 respectively. The extractors can function on any SPARQL endpoint irrespective of its (i) version, (ii) support for different SPARQL keywords<sup>39</sup>, and (iii) the size of the underlying RDF graphs. The Extraction Algorithm can generally process most SPARQL endpoints within 4–8 hours. A visualization of a portion of the extracted  $\mathcal{G}_{LSLOD}$  schema graph is shown in Fig. 2b. A snippet of the extracted information for a given source is shown in Box 2.

Through an empirical analysis, we have found some overlap between the different schema elements, based on the way linked data publishers model the underlying data (i.e., some object properties in a given source are used as data properties in another LOD source). We found standard datatypes defined under the XML Schema Definition (XSD)<sup>77</sup> and RDFS<sup>36</sup> such as `xsd:string`, `xsd:integer`, `xsd:float`, `xsd:dateTime`, `xsd:boolean`, `rdfs:Literal`. The analysis also found an ontology class `UO:0000034` (Units Ontology<sup>78</sup> Week class) used as a datatype. In some cases, we also found that RDF graphs may exhibit semantic mismatch where instances are aligned to classes from an exhaustive ontology, such as ChEBI<sup>6</sup> or NCIT<sup>15</sup> that have more than 50,000 classes, using the `rdf:type` property. *Semantic Mismatch* in LOD is a form of semantic heterogeneity. Since `rdf:type` is a predicate in almost all the SPARQL query templates used by the Extraction Algorithm, semantic mismatch significantly increases the number of classes in the  $\mathcal{G}_{T-Box}$  and consequently the number of queries that need to be formulated for each class significantly increases.

**Vocabulary reuse across biomedical linked open data sources.** We analyzed the 70,592 total schema elements (classes, object properties, and data properties) extracted from the LOD sources in  $\mathcal{G}_{LSLOD}$  for vocabulary



**Fig. 4** Vocabulary reuse across biomedical linked open data sources (RDF graphs): **(a)** A histogram of schema elements that are shared across multiple LOD Sources. It can be seen that most LOD Sources use a unique schema or vocabulary whose elements are not reused across any other data source. However, there are at least 100 schema elements that are shared by >10 LOD Sources. **(b)** The different vocabularies (published at LOV repository) and ontologies (BioPortal and other OWL ontologies) whose elements are reused in the schemas of LOD sources. Several LOD sources either use unpublished schemas or use elements from common vocabularies (e.g. DCTerms and SKOS).

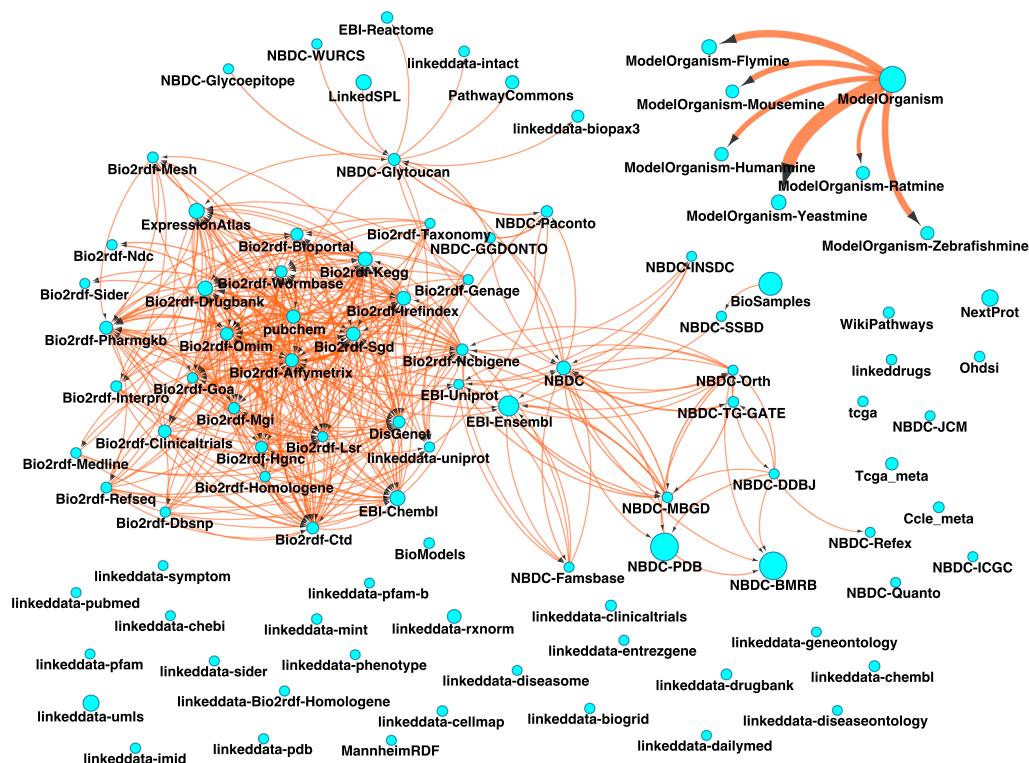
reuse. In this research, we looked only at reuse of schema elements from vocabularies published at the Linked Open Vocabularies (LOV) repository<sup>45</sup> and ontologies that are either published at the BioPortal repository<sup>46</sup> or on the Web as OWL files available for download (and have .owl in the namespace prefix, for easy search).

Figure 4 shows the results of the analyses. Figure 4a depicts a histogram of schema elements that are shared across  $X$  number of LOD sources. It can be observed that most LOD sources are composed of unique schema elements that are present only in that particular source (i.e., the first histogram bar from the left). However, there are at least >100 schema elements that are shared between >10 LOD sources (i.e., the last four histograms from the right).

Terms are reused in biomedical LOD graphs from 34 BioPortal ontologies, 7 other biomedical ontologies, available on the Web but not stored in the BioPortal repository, and 43 LOV vocabularies (Fig. 4b). Some of the reused ontologies and vocabularies are very popular in biomedicine (e.g., ChEBI<sup>6</sup>, SNOMED CT<sup>62</sup>, NCIT — National Cancer Institute Thesaurus<sup>15</sup>, GO — Gene Ontology<sup>19</sup>) and in the Semantic Web community (e.g., DCTerms Metadata Description Model<sup>79</sup>, RDFS<sup>36</sup>, SKOS — Simple Knowledge Organization System<sup>64</sup>, Data Catalog Vocabulary<sup>80</sup>) respectively.

However, most of these biomedical ontologies are used only in the schemas of a few LOD sources ( $X$ -axis), and moreover only a few elements from these ontologies are reused ( $Y$ -axis). For example, the Anatomical Therapeutic Chemical Classification terminology<sup>81</sup> — popularly used to classify drugs according to their mechanism of action — is used in the schema of only one LOD source. It can be seen that several LOD sources use unpublished schemas (green diamonds), hence, a large number of unique schema elements are present in only one source in Fig. 4a.

Several LOD sources do reuse schema elements from a set of widely used vocabularies ( $\approx 100$  as per Fig. 4a). However, the elements of these vocabularies are not useful for data integration and integrated querying from a biomedicine perspective. It should be noted that several elements from the combination of the Semantic Science Ontology and Chemical Information Ontology (SIO-CHEMINF)<sup>47,82</sup> are reused across several LOD sources and may be a popular resource for the emergence of a common data model for integrated querying of these sources. A large number of these LOD sources however only use one particular element from SIO-CHEMINF ontology.



**Fig. 5** Inter-linking and intra-linking across biomedical linked open data sources (RDF graphs): Each unique RDF Graph retrieved by the Extraction Algorithm is shown as a node in this network. The size of the node indicates the number of unique object properties used to link entities in different classes from the same source. The size of the edge indicates the number of unique object properties used to link entities in different classes in different sources.

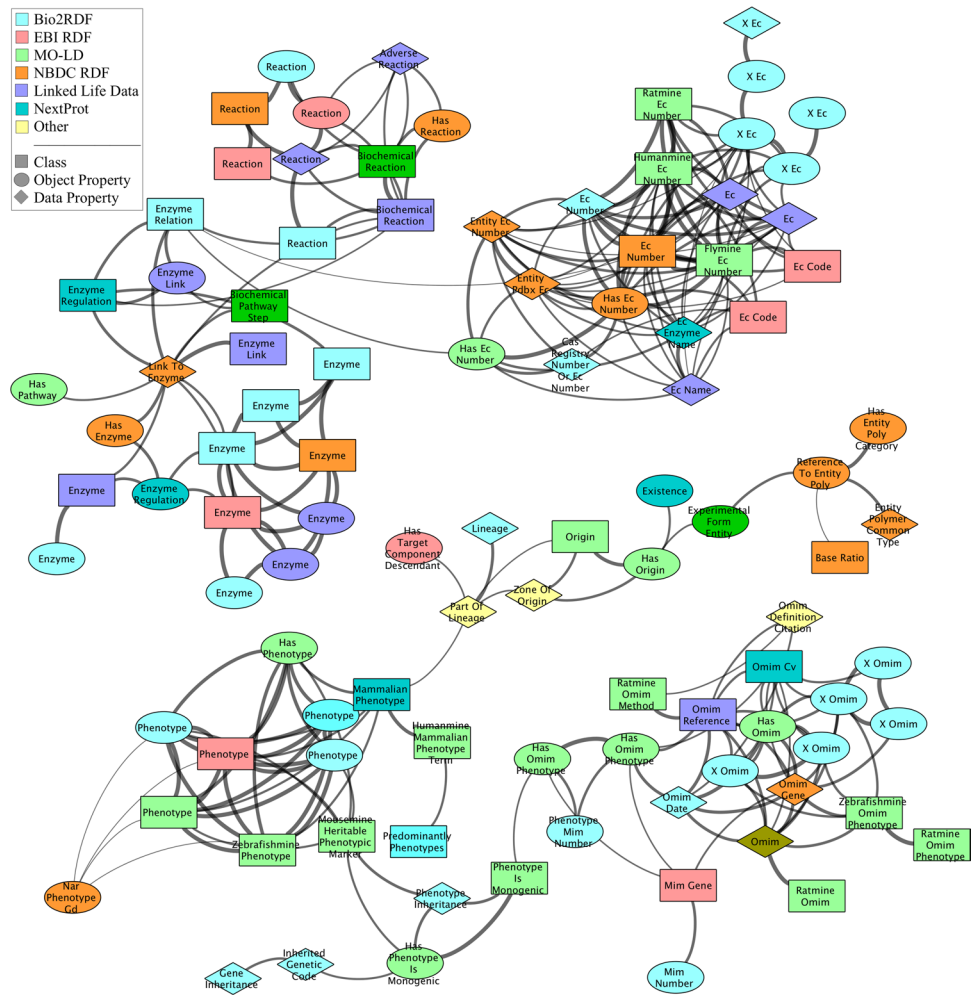
Using Eq. 1, vocabulary reuse was determined to be 19.86%, which may indicate favorable reuse across LOD sources in comparison to ontologies. However, it should be noted that this statistic may be driven by the reuse of 30,251 ChEBI classes in a few sources. We also determined several cases of *semantic mismatch* where classes from popular biomedical ontologies in the BioPortal repository (e.g., SNOMED CT, NCIT, GO) are mapped to instances in these LOD sources.

Figure 5 shows the different RDF graphs extracted as nodes in a network. The size of the node indicates the number of unique object properties used to link entities in different classes from the same source (intra-linking). The size of the edge connecting two different nodes indicates the number of unique object properties used to link entities in different classes in different sources or RDF Graphs (represented using the corresponding node). This network diagram can be perceived to be similar to the LSLOD cloud diagram<sup>23</sup> (available at <http://lod-cloud.net>).

It can be observed that the LSLOD cloud is not actually as densely connected as is often indicated through the LSLOD cloud diagram. Several sources, which may exhibit a higher degree of intra-linking are not inter-linked with other RDF graphs and exist as stand-alone data sources converted using RDF. It can also be observed that the PDB RDF graph, exposed by the NBDC project, and the UniProt RDF graph, published by the EBI-RDF consortium, show a lack of connectivity. While the underlying data sources are similar and consist of information on proteins, the lack of connectivity can be explained if the NBDC project uses a different representation scheme for URIs used to denote protein entities as compared to the EBI-RDF consortium. The highly appreciated efforts by the Bio2RDF<sup>20</sup> and EBI-RDF<sup>60</sup> consortia can be seen in the dense inter-linking between the consortia-published sources.

**Communities of related content in the LSLOD cloud.** While the low level of vocabulary reuse across LOD sources present a bleak picture of the Semantic Web vision toward facilitating Web-scale computation and integrated querying, it should be noted that there is significant similarity of content across these sources.

Using the pre-trained word embedding vectors and IDF statistics, we generated 100-dimensional URI embedding vectors for different schema elements URIs in the  $\mathcal{G}_{LSLOD}$  schema graph. Whereas the word tokens for most URI labels were found in the vocabulary, 3,340 word tokens, such as ‘differrn’, ‘heteronucl’, and ‘neoplas’ were not found either due to typos or due to minimal occurrence frequency in publications. We represented these words using the default embedding vector and IDF statistic. We computed cosine similarity scores to detect additional mappings between different schema elements. For example, the `drugbank:MolecularWeight` class has higher cosine similarity scores to related schema elements, such as `kegg:mol_weight`, `biopax:molecularWeight`, and `MOLD:hasMolecularWeight` data properties, as well as the `CHEMINF:000198` (MOLECULAR WEIGHT CALCULATED BY PIPELINE PILOT) class.



**Fig. 6** Communities of similar content across biomedical linked open data: Two communities composed of schema elements related to ENZYME and PHENOTYPE concepts are visualized in a greater detail. The nodes indicate a distinct schema element (distinct URI), and are colored according to the LOD project they are present in. The shape of the nodes indicates the type of schema element (class, object property, data property). The width of an edge between two nodes is proportional to the cosine similarity score between the URI embeddings of the corresponding schema element. Smaller sub-communities related to REACTION, *Enzyme Commission Number* and *OMIM Phenotype Identifier* (Online Mendelian Inheritance in Man) are also observed.

A similarity network was created with 12,613 schema elements as nodes in this network. This number of nodes is less than the total number of schema elements in the  $\mathcal{G}_{LSLOD}$  schema graph (70,587). We excluded domain-specific classes that were reused from ChEBI and other ontologies through semantic mismatch. We also excluded URIs for which labels could not be extracted.

Edges were created between these nodes using the cosine similarity scores between the corresponding schema elements. There were 169,005 pairs of schema element URIs that satisfied the first criterion of having a cosine similarity value above the threshold of 0.75. This threshold was predetermined after an exploratory analysis of pairs of schema elements that satisfied this criterion. While it would make sense to have a higher threshold ( $\approx 0.95$ ) for a higher degree of certainty on similarity, having a threshold of 0.75 allowed the inclusion of pairs, such as `obo:regulates`  $\leftrightarrow$  `drugbank:mechanism-of-action:0.75`, `chembl:-mechanismDescription`  $\leftrightarrow$  `drugbank:mechanism-of-action:0.85`, `faldo:location`  $\leftrightarrow$  `drugbank:cellular-location:0.89`. After applying the second criterion of one-to-one mapping between two schema elements for a particular combination of LOD sources, the network had 22,855 pairs of schema-element URIs for which edges were created.

We used the Louvain method for detecting communities in the weighted similarity network. The method terminated after achieving a maximum modularity statistic (i.e., the density of edges inside communities compared to edges between communities) of 0.817296. We detected 6,641 communities. While most of these communities had minimal membership (i.e.,  $\leq 10$  schema elements were members), there were 51 communities with more than 10 schema elements as members. On visual inspection using the Cytoscape network visualization platform, we determined that these communities primarily consisted of schema elements related to a few key biomedical concepts or relation types. These communities are partially visualized in Fig. 3, with two communities composed

of schema elements related to ENZYME and PHENOTYPE concepts visualized in a greater detail in Fig. 6. ENZYME entities are specialized proteins which act as biological catalysts to accelerate biochemical reactions, whereas PHENOTYPE entities are composite observable characteristics or traits of an organism.

In Fig. 6, each node is a distinct schema element URI (i.e., class, object property or data property) observed in the  $\mathcal{G}_{LSLOD}$  LSLOD schema graph. These nodes are colored based on the LOD project they are primarily present in (e.g., Bio2RDF, EBI-RDF). It should be noted that each LOD project (e.g., Bio2RDF) may publish multiple sources as distinct RDF graphs (e.g., KEGG<sup>9</sup>, DrugBank<sup>11</sup>). The shape of the nodes is reminiscent of the type of schema element — classes are represented as rectangular nodes, object properties as elliptical nodes, and data properties as diamond nodes. The nodes are labeled using the extracted labels, but the underlying URIs may be completely different. The width of the edge connecting different nodes is proportional to the cosine similarity scores between the URI embeddings of the connected schema elements.

It can easily be seen that different schema element URIs may be used across different sources (even if the sources are published by the same LOD project) to model and represent similar content and information (e.g., information on enzymes or phenotypes). It also showcases that data and knowledge relevant to a particular biomedical entity may be scattered across multiple sources (i.e., knowledge on a particular Enzyme entity can be retrieved from Bio2RDF, EBI-RDF and NextProt, and in many cases this knowledge may be novel). Other communities composed of schema elements focused on other key biomedical concept or relation types are also observed (e.g., DRUG, PUBLICATION, CHROMOSOME). The primary composition of a few of these larger communities (>30 members) is listed in Table 2. It should be noted that these communities may be composed of multiple sub-communities. For example, a sub-community of schema elements related to the REACTION concept is linked to the main cluster of nodes related to the ENZYME concept, since ENZYME entities are often associated with REACTION entities.

Smaller communities (<10 members) also consist of specific schema elements in different sources that are similar to each other but that are in label mismatch through use of different URIs (e.g., snRNA is represented using

MOLD:yeastmine\_SnoRNAGene, MOLD:humanmine\_SnoRNA, EBIRDF:ensembl/snoRNA, BIO2RDF:ncbigene\_vocabulary/SnoRNA-Gene, NBDC:mbgd#snoRNA). Knowledge of such communities will aid biomedical researchers to determine which sources in the LSLOD cloud contain information relevant to their domain of interest (e.g., research on small nucleolar RNAs snRNA, a class of small RNA molecules that primarily guide chemical modifications).

## Discussion

Semantic heterogeneity across different data and knowledge sources, either in the use of varying schemas or entity notations, severely hinders the retrieval and integration of relevant data and knowledge for use in downstream biomedical applications (e.g., drug discovery, pharmacovigilance, question answering). Enhancing the quality and decreasing the heterogeneity of biomedical data and knowledge sources will lead to enhanced inter-disciplinary biomedical research, which will lead to developing better clinical outcomes and increasing our knowledge on biological mechanisms.

So far, significant resources have been invested to represent, publish, and link biomedical data and knowledge on the Web using Semantic Web technologies to create the Life Sciences Linked Open Data (LSLOD) cloud. However, there are very few biomedical applications that query and use *multiple* LOD sources and biomedical ontologies directly on the Web to generate new insights, or to discover novel implicit associations *serendipitously*. This is, in part, due to the semantic heterogeneity emanating from underlying sources, which has also penetrated the LSLOD cloud in varying forms. Some of the primary benefits of the LSLOD cloud — querying, retrieval, and integration of heterogeneous data and knowledge from multiple sources on the Web seamlessly — cannot be availed.

Linked data publishers are strongly encouraged to reuse existing models and vocabularies and to provide links to other sources in the LSLOD cloud<sup>21,41</sup>. Similarly, biomedical ontology developers are strongly encouraged to reuse equivalent classes existing in other ontologies, while building new ontologies<sup>83,84</sup>. There are several benefits if biomedical ontologies and LOD sources reuse content from existing sources — (i) reduction in data and knowledge engineering and publishing costs, (ii) decrease in semantic heterogeneity and increase in semantic interoperability across datasets, and (iii) ease of querying, retrieval, and integration of data from multiple sources simultaneously through existing query federation methods<sup>48</sup>. On the other hand, lack of reuse will manifest as increased semantic heterogeneity across different LSLOD sources.

It can be asserted through the findings presented in this research that the Life Sciences “Linked” Open Data cloud is not actually as densely connected as is often indicated through the ubiquitously-referenced LSLOD cloud diagram. Several sources, which may exhibit a higher degree of intra-linking are not inter-linked with other RDF graphs and exist as stand-alone data sources converted using RDF (Fig. 5). Similarly, Fig. 4 shows that several biomedical LOD sources use unpublished schemas. Biomedical LOD sources do reuse schema elements from a set of popular vocabularies. However, these vocabularies originate from the Semantic Web community (e.g., those that are mainly used to refer to data elements and attributes) rather than from the biomedical community (e.g., those that are used to provide normalization schemes for biomedical entities). Hence, the elements of these vocabularies are not useful for data integration and integrated querying from a biomedical perspective (e.g., for retrieving all drug–protein target interactions).

**Similar content across the LSLOD cloud.** As seen in Fig. 6, several online sources have data and knowledge pertaining to ENZYME or PHENOTYPE entities but use completely different schema elements. Communities of similar content depicted in Fig. 6, while showcasing the different schema element URIs to represent similar content, may also provide an idea on the different equivalence attribute URIs used to link entities across

Com. Id.	Total Elements	Key Biomedical Concepts
5023	549	<b>Clinical</b> (Age, Treatment, Summary, Dosage, Diagnosis)
		<b>Cancer</b> (Stage, Tumor, Cell, Sample, Rate)
		<b>Genomics</b> (Experiment, Target, Function, Interaction, Correlation, Assembly, Association, Expression, Genetic)
3098	281	<b>Protein and Genomic Sequences</b> (Sequence, Position, Start, Length, Protein, End, Repeat, Gene, Match, Begin, Translation, Exact, Exon, Frame, Distance)
2842	271	<b>Clinical</b> (Date, Condition, Entry, Process, Person, Diagnosis, Participant, Outcome, Indication, Therapy, Evidence, Access)
4346	260	<b>Data</b> (Code, Reference, Database, Sample, Data, Ontology, Semantic, Experiment, Dataset, Assembly, Registry)
2879	231	<b>Clinical</b> (Clinical, Group, ICD)
		<b>Genetic</b> (Gene, Entity, Ncbigene, Species, Atom, Ref)
3945	228	<b>Biomedical Relations</b> (Disease, Relationship, Expression, Gene, Level, Sample, Drug, Interaction, Component)
2515	211	<b>Protein</b> (Protein, Site, Region, Molecule, Domain, Binding, Interaction, Active, Atom, Canonical, Primary, Structure)
63676	210	<b>Literature</b> (Description, Author, Entity, Details, Note, Article, Issue, List, Type, Study, Entry, Document, Data, Address)
1417	192	<b>Identifiers</b> (Uniprot, Identifier, Accession, Ensembl, Taxonomy, PDB, Genbank, Gene, MGI, Database, Id, Refseq, NCBI, CDS)
1632	171	<b>Identifiers</b> (Id, Taxon, Mesh, Entity, Type)
		<b>Genomics</b> (Platform, Entry, Sample, Genus, Array, Gene, List)
3249	171	<b>Genotype-Phenotype Associations</b> (Gene, Symbol, Association, Product, Variant, Tag, Locus, Phenotype, Allele, Disease)
5652	158	<b>Genomics</b> (RNA, Probe, Strand, Direction, Region, DNA)
1446	132	<b>Software</b> (Version, Software, File, Format, Image, Email, Scale)
2535	126	<b>Drug</b> (Drug, Chemical, Action, Ingredient, Bond, Activity)
2421	110	<b>Chromosome</b> (Location, Chromosome, Entity, Cellular)
2317	105	<b>Pathway</b> (Pathway, Molecular, Formula)
		<b>Unit</b> (Weight, Body, Units, Volume, Unit)
1806	101	<b>Source</b> (Source, Natural, Vector, Derived, Tissue)
2672	99	<b>Literature</b> (Count, Text, Comment, Editor, List, Data, Notes)
6316	97	<b>Ontology</b> (Annotation, Ontology, Type, Structure)
5969	83	<b>Social</b> (Status, Gender, Authority, Interaction)
2429	79	<b>mRNA</b> (Transcript, mRNA, Signal, Frequency, Expression)
1868	77	<b>Literature</b> (Citation, Publication, Journal, Literature)
1349	67	<b>Assay</b> (Method, Assay, Type, Measurement)
6241	62	<b>Genetic</b> (Strain, Allele, Mutation, Genotype, SNP)
6052*	50	<b>Enzyme</b> (EC, Enzyme, Reaction)
3524*	50	<b>Phenotype</b> (Phenotype, Omim)
2490	50	<b>Organization</b> (Component, Complex, Assembly, Entity)
1849	49	<b>Interaction</b> (Interaction, Experiment, Strength)
2897	49	<b>Cell Line</b> (Cell, Line, Tissue, Type, Node)
6283	39	<b>Organism</b> (Organism, Host, Entity)
2911	38	<b>Ontology</b> (Synonym, Term, Ontology)
1787	32	<b>Evolution</b> (Member, Family, Orthologue, Homologue, Evidence)
5695	31	<b>Literature</b> (Pubmed, Keyword)

**Table 2.** Primary composition of communities in terms of key biomedical concepts. (\*) marked communities are visualized in Fig. 6.

different sources to a specific entity coding scheme. For example, the ENZYME entities in the different sources of the LSLOD cloud are linked to the Enzyme Commission number (a numerical classification scheme for enzymes) using different attribute URIs with labels such as Ec CODE, Ec NUMBER, X Ec, and HAS Ec NUMBER. Similarly, PHENOTYPE entities are linked to the OMIM phenotype identifiers (i.e., Online Mendelian Inheritance in Man is a knowledge base of human genes and genetic phenotypes<sup>85</sup>) using different attribute URIs with labels such as OMIM REFERENCE, X OMIM, and HAS OMIM PHENOTYPE. Communities that are composed entirely of different attribute URIs referring entities in different sources to a uniform coding scheme are also observed (e.g., ENSEMBL and HGNC Gene Coding scheme attributes). The knowledge on the different equivalence attribute URIs could guide query federation methods to perform entity reconciliation after the retrieval of information from multiple sources.

Through a manual inspection of such “identifier” communities (e.g., communities 1417 and 1632 in Table 2), we were able to detect different URI representations for similar entities (classes and instances). A few examples of different URI representations are shown in Table 3 for chemical entities of ChEBI<sup>6</sup>, PubChem<sup>5</sup> and MeSH<sup>86</sup>



Entity Origin	URI Representation	Sources
ChEBI	<a href="http://purl.obolibrary.org/obo/CHEBI/*">http://purl.obolibrary.org/obo/CHEBI/*</a>	LinkedSPL, PubChem, NBDC ExpressionAtlas, BioSamples
	<a href="http://identifiers.org/chebi/CHEBI:">http://identifiers.org/chebi/CHEBI:</a>	BioModels, PDB, WikiPathways Bio2RDF, PathwayCommons
	<a href="http://identifiers.org/obo.chebi/CHEBI:">http://identifiers.org/obo.chebi/CHEBI:</a>	BioModels, BioPax, WikiPathways
	<a href="http://www.ebi.ac.uk/chebi/">http://www.ebi.ac.uk/chebi/</a>	PDB, ChEMBL
	<a href="http://bio2rdf.org/chebi:">http://bio2rdf.org/chebi:</a>	Bio2RDF
PubChem	<a href="http://identifiers.org/pubchem.compound/">http://identifiers.org/pubchem.compound/</a>	Bio2RDF, WikiPathways, NBDC
Compound	<a href="http://rdf.ncbi.nlm.nih.gov/pubchem/compound/*">http://rdf.ncbi.nlm.nih.gov/pubchem/compound/*</a>	WikiPathways, PubChem
	<a href="http://bio2rdf.org/pubchem.compound:">http://bio2rdf.org/pubchem.compound:</a>	Bio2RDF
	<a href="http://pubchem.ncbi.nlm.nih.gov/compound/">http://pubchem.ncbi.nlm.nih.gov/compound/</a>	NBDC
MeSH	<a href="http://purl.bioontology.org/ontology/MESH/*">http://purl.bioontology.org/ontology/MESH/*</a>	NBDC, PDB, LinkedSPL
	<a href="http://id.nlm.nih.gov/mesh/YEAR/*">http://id.nlm.nih.gov/mesh/YEAR/*</a>	MeSH, DisGenet, PDB
	<a href="http://identifiers.org/mesh/">http://identifiers.org/mesh/</a>	Bio2RDF, DisGenet, PDB
	<a href="http://www.ncbi.nlm.nih.gov/mesh/">http://www.ncbi.nlm.nih.gov/mesh/</a>	BioSamples
	<a href="http://rdf.imim.es/rh-mesh.owl#">http://rdf.imim.es/rh-mesh.owl#</a>	DisGenet
	<a href="http://bio2rdf.org/mesh_resource:">http://bio2rdf.org/mesh_resource:</a>	Bio2RDF
	<a href="http://bio2rdf.org/mesh:">http://bio2rdf.org/mesh:</a>	Bio2RDF

**Table 3.** Different kinds of URI representations for chemical compounds observed in LOD Sources. (\*) marks the recommended representation(s). **Acronyms:** (NBDC) National Bioscience Database Center, (PDB) Protein Data Bank, (MeSH) Medical Subject Headings.

sources. Most LOD querying architectures rely on the existence of exact URIs to query multiple sources simultaneously (using data warehousing, link traversal or query federation methods)<sup>13,17,40</sup>. However, since these URIs are essentially different, querying architectures will not be able to integrate data and knowledge from different sources (e.g., protein–ligand binding information from PDB and biological assay experimental data from PubChem). These examples qualify as “intent for reuse” and add to the challenges of semantic heterogeneity across the LSLOD cloud.

**Intent for reuse across the LSLOD cloud.** A few LOD sources do reuse content from existing biomedical ontologies, but in many cases the reuse is limited to very few classes or involve incorrect representations. Previously, we have documented similar cases to be “intent for reuse” on the part of ontology engineers in the biomedical domains<sup>48</sup>. Table 3 documents similar cases of intent for reuse that were observed across the biomedical LOD sources. Essentially, these URI patterns generally have the same identifier and source ontology or vocabulary, but are reused from different versions of the source ontology or vocabulary, or represented using different notations or namespaces. These patterns cannot be considered as actual reuse, as these are different, and often incorrect, representations for the same terms, and no explicit mappings are found. Hence, the advantages of reuse cannot be experienced. By using the correct representations, the semantic heterogeneity (i.e., content overlap) can be reduced substantially. Ideally, the representation of the URI used at the “authoritative source” (i.e., the source where the respective semantic resource is published and regularly maintained) is the standard representation that should be consistently reused across other LSLOD sources. Determining the “authoritative source” is out of the scope for this meta-analysis, but is a field of active research<sup>50</sup>. The recommended representations in Table 3 are based off the “authoritative source”. Ontology engineers and linked data publishers may lack sufficient guidelines and tools to identify authoritative sources for any semantic resource and use the recommended representation while reusing or referring to that semantic resource.

Intent for reuse and semantic heterogeneity severely hinders most query federation and link traversal methods for heterogeneous data and knowledge integration across the LSLOD cloud. For a more concrete example, consider the user query “Retrieve Assays which identify potential Chemopreventive Agents that have Molecular Weight < 300 g/mol and target Estrogen Receptor present in Human.” This query requires the retrieval of knowledge from the KEGG knowledge base<sup>9</sup> on all the chemicals that target ESTROGEN RECEPTOR protein, and then the retrieval of assay data on these chemicals from the ChEMBL chemical bioactivity repository<sup>59</sup>. The KEGG knowledge base and the ChEMBL data repository are both published on the LSLOD cloud<sup>20,60</sup>, and have chemical URIs that have `x-ref` links to chemical URIs in the ChEBI ontology of biological chemicals<sup>6</sup>. However, since Bio2RDF KEGG uses the URI representation scheme [http://bio2rdf.org/chebi:\\*](http://bio2rdf.org/chebi:*) and ChEMBL uses the URI representation scheme [http://purl.obolibrary.org/obo/CHEBI/\\*](http://purl.obolibrary.org/obo/CHEBI/*) for the same identifiers in the ChEBI ontology (Table 3), it is impossible to formulate a federated SPARQL query that can query KEGG and ChEMBL, simultaneously and retrieve the relevant results in an integrated fashion.

Increasing the reuse of schema elements from existing vocabularies and ontologies may decrease semantic heterogeneity and enable the integrated querying of multiple data and knowledge sources across the LSLOD cloud using the Semantic Web technologies. In cases where linked data publishers have the need to use a custom vocabulary or ontology, the overarching advice would be to publish this vocabulary or ontology to a popular repository, so other publishers can use these resources. The challenges for data and knowledge integration due

to semantic heterogeneity across fragmented and incompatible data sources will also prevail as the wider data science community moves toward the adoption of the FAIR (Findable, Accessible, Interoperable, Reusable) principles<sup>87</sup> for publishing and reusing open datasets on the Web.

Over the years, the Semantic Web community has recommended several best practices and standards on how to reuse schema elements from existing vocabularies and ontologies when publishing LOD sources<sup>30,41,50,88</sup>. However, centralized top-down approaches may be needed to evaluate and enforce the use of these standards and existing ontologies and vocabularies (e.g., SKOS<sup>64</sup>, Schema.org<sup>65</sup>, Semantic Science Ontology<sup>82</sup>) when a new source is published in the LOD cloud. Such centralized approaches are already being developed to monitor the accessibility and availability of the different sources in the LOD cloud<sup>30,33,89,90</sup>, and can be extended to avoid the proliferation of incompatible and redundant LOD sources and reduce our reliance on reuse standards and best practices that would emerge through an evolutionary process.

**Limitations and future work.** The use of word embeddings in the community detection algorithm enables the identification of mappings between similar LOD schema elements (e.g., `kegg:therapeutic-category` ↔ `drugbank:Drug-Classification-Category`: 0.93, `obo:regulates` ↔ `drugbank:mechanism-of-action`: 0.75, `faldo:location` ↔ `drugbank:cellular-location`: 0.89). These mappings may provide suggestions to linked data publishers and consumers on relevant information in different sources. However, many of these mappings may not be accurate due to reliance on the word embeddings, as well as the 0.75 threshold on the cosine similarity scores. For example, as seen in Fig. 6, while the REACTION concept in the ENZYME community symbolizes biochemical reactions where enzyme entities may act as catalyzing agents, an error with the community detection algorithm also indicates that ADVERSE REACTION (clinical concept observed in patients treated with a set of drugs) may be a part of this community. This is due to the fact that the singular schema element pertaining to ADVERSE REACTION from an RDF graph in the Linked Life Data project<sup>91</sup> is linked to schema elements pertaining to BIOCHEMICAL REACTION (i.e., the edges have a cosine similarity score of more than 0.75).

In this research, we have limited this meta-analysis to only include the BioPortal repository<sup>46</sup> and Linked Open Vocabularies (LOV) catalog<sup>45</sup> as background sources of biomedical ontologies and LOD vocabularies. However, several other repositories may contain ontologies and vocabularies which are useful toward this meta-analysis (e.g., Ontology Lookup Service – OLS<sup>92</sup>). There were two underlying assumptions behind this decision: *i*) There may be very few biomedical ontologies and vocabularies that are not present in BioPortal and LOV catalog, but are exposed in another public source (e.g., As of October 2020, OLS has 252 ontologies whereas BioPortal has 896 ontologies, and only around 40 ontologies in OLS are not present in BioPortal. Out of these 40 ontologies, it should also be noted that 30 ontologies are “organism” ontologies which may have domain-specific concepts not present in the LSLOD cloud), and *ii*) the LSLOD cloud diagram at <https://lod-cloud.net> indicates that LOD sources are more intricately linked to biomedical ontologies in BioPortal repository. This meta-analysis can be extended in the future to include other additional sources of background knowledge.

While we have tried to conduct the meta-analysis by classifying the LSLOD sources according to the project that publishes them (e.g., Bio2RDF, EBI, NBDC), a further analysis can be conducted in conjunction with author citation networks to discern the common publishers across different LSLOD sources (e.g., publishers of Bio2RDF<sup>20</sup> sources may be involved with the publishing efforts of WikiPathways<sup>61</sup> or PubChem<sup>5</sup> RDF graphs). This would enable us to determine the true independent publishers of LSLOD sources. Ideally, data publishing, curation, reusability, and governance should be a community-driven process, with monitors to ensure that the best practices and standards are followed by the different RDF sources.

For this research, the meta-analysis was conducted over the April 2018 snapshot of the LSLOD cloud, as well as using the biomedical ontologies and RDFS vocabularies from that time period. However, the LSLOD cloud is constantly evolving and the biomedical data and knowledge sources are published, updated, or deprecated on a regular basis. In the future, the methods presented in this research can be embedded within a Web application which automatically extracts schema elements from any novel biomedical data or knowledge source, which is registered by linked data publishers, and analyzes these elements in the context of the current LSLOD cloud and biomedical resources stored in public catalogs (e.g., BioPortal<sup>46</sup> and the LOV catalog<sup>45</sup>). A formatted report with analyses results and recommendations can be issued to the data publishers empowering them to improve reuse in their source. This Web application, while outside the scope of the current meta-analysis research, can also continually monitor the LSLOD cloud to generate insights on the evolution, availability, quality, and semantic heterogeneity, of the different sources in the LSLOD cloud on a regular basis.

## Conclusion

Semantic Web and linked data technologies garner significant interest and investment from the biomedical research community toward tackling the diverse challenges of heterogeneous biomedical data and knowledge integration. In this paper, we conduct an empirical meta-analysis to demonstrate that there is considerable heterogeneity and discrepancy in the quality of Life Sciences Linked Open Data (LSLOD) sources on the Web, which are published using the above technologies. We discuss on whether the LSLOD cloud can truly be considered “linked” due to the heterogeneous schemas, varying entity notations, the lack of mappings between similar entities, and the lack of reuse of common vocabularies. While increasing reuse across biomedical ontologies and LOD sources is definitely considered to be an attractive alternative to decrease semantic heterogeneity, it will take a while for biomedical ontology developers and linked data publishers to realize the incentives of reuse and actually embrace the guidelines and best practices while publishing data and knowledge. Actual reuse through the use of correct URI representations in biomedical ontologies and LOD sources, rather than “intent to reuse”, can decrease semantic heterogeneity substantially itself. The findings from our meta-analysis across the LSLOD cloud, as well as the resources such as the LSLOD schema graph

made available through this research, will lead to the development of better tools and methods to increase vocabulary reuse and enable efficient querying and integration of multiple heterogeneous biomedical sources on the Web.

### Data availability

The LSLOD Schema Graph extracted from the different SPARQL endpoints of the biomedical LOD graphs in the LSLOD cloud is made available in the JavaScript Object Notation (JSON) format serialized as a Pickle file in the figshare repository (<https://doi.org/10.6084/m9.figshare.12402425.v2>) under a CC BY 4.0 license<sup>93</sup>. The project also contains tab-separated values files consisting of information about extracted classes, object properties, data properties, datatypes, as well as of statistics on the LOD graphs from the LSLOD cloud used in this research.

### Code availability

The different visualizations and observations presented in this research are also made available at the project website (<http://onto-apps.stanford.edu/lslodminer>). The scripts for the extraction algorithm and the community detection algorithm used in this research are made available on GitHub (<https://github.com/maulikkamdar/LSLODQuery>).

Received: 4 June 2020; Accepted: 4 December 2020;

Published online: 21 January 2021

### References

1. Johnson, A. E. *et al.* MIMIC-III, a freely accessible critical care database. *Scientific data* **3**, 160035 (2016).
2. Clark, K. *et al.* The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository. *Journal of digital imaging* **26**, 1045–1057 (2013).
3. Weinstein, J. N. *et al.* The cancer genome atlas pan-cancer analysis project. *Nature genetics* **45**, 1113 (2013).
4. US National Libraries of Medicine. PubMed. <https://www.ncbi.nlm.nih.gov/pubmed/> (2018). [Online; accessed 19-July-2018].
5. Fu, G. *et al.* PubchemRDF: towards the semantic annotation of pubchem compound and substance databases. *Journal of cheminformatics* **7**, 34 (2015).
6. Hastings, J. *et al.* The ChEBI reference database and ontology for biologically relevant chemistry: enhancements for 2013. *Nucleic acids research* **41**, D456–D463, <https://doi.org/10.1093/nar/gks1146> (2013).
7. Consortium, U. *et al.* The universal protein resource (UniProt). *Nucleic acids research* **36**, D190–D195 (2008).
8. Croft, D. *et al.* The Reactome pathway knowledgebase. *Nucleic acids research* **42**, D472–D477 (2014).
9. Kanehisa, M. & Goto, S. KEGG: Kyoto Encyclopedia of Genes and Genomes. *Nucleic acids research* **28**, 27–30 (2000).
10. Mungall, C. J. *et al.* The Monarch initiative: an integrative data and analytic platform connecting phenotypes to genotypes across species. *Nucleic acids research* **45**, D712–D722 (2017).
11. Wishart, D. S. *et al.* DrugBank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic acids research* **34**, D668–D672 (2006).
12. Bodenreider, O. Biomedical ontologies in action: role in knowledge management, data integration and decision support. *Yearbook of medical informatics* **67** (2008).
13. Kamdar, M. R., Fernández, J. D., Polleres, A., Tudorache, T. & Musen, M. A. Enabling web-scale data integration in biomedicine through linked open data. *NPJ digital medicine* **2**, 1–14 (2019).
14. Williams, A. J. *et al.* Open PHACTS: semantic interoperability for drug discovery. *Drug discovery today* **17**, 1188–1198, <https://doi.org/10.1016/j.drudis.2012.05.016> (2012).
15. Sioutos, N. *et al.* NCI Thesaurus: a semantic model integrating cancer-related clinical and molecular information. *Journal of biomedical informatics* **40**, 30–43, <https://doi.org/10.1016/j.jbi.2006.02.013> (2007).
16. Jentzsch, A. *et al.* Linking Open Drug Data. In *I-SEMANTICS* (2009).
17. Saleem, M. *et al.* Big linked cancer data: Integrating linked TCGA and PubMed. *Web Semantics: Science, Services and Agents on the World Wide Web* **27**, 34–41 (2014).
18. Bodenreider, O. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research* **32**, D267–D270 (2004).
19. Ashburner, M. *et al.* Gene ontology: tool for the unification of biology. *Nature genetics* **25**, 25 (2000).
20. Callahan, A., Cruz-Toledo, J., Ansell, P. & Dumontier, M. Bio2RDF release 2: Improved coverage, interoperability and provenance of life science linked data. In *The Semantic Web: Semantics and Big Data*, 200–212 [https://doi.org/10.1007/978-3-642-38288-8\\_14](https://doi.org/10.1007/978-3-642-38288-8_14) (Springer, 2013).
21. Bizer, C., Heath, T. & Berners-Lee, T. Linked data—the story so far. *Semantic Services, Interoperability and Web Applications: Emerging Concepts* 205–227 (2009).
22. Berners-Lee, T., Hendler, J. & Lassila, O. *et al.* The semantic web. *Scientific american* **284**, 28–37 (2001).
23. Abele, A., McCrae, J. P., Buitelaar, P., Jentzsch, A. & Cyganiak, R. Linking open data cloud diagram (2017) (2017).
24. Ramaswami, P. A remedy for your health-related questions: health info in the knowledge graph. *Google Official Blog* 2018 (2015).
25. AOCNP, D. Watson will see you now: a supercomputer to help clinicians make informed treatment decisions. *Clinical journal of oncology nursing* **19**, 31 (2015).
26. Kamdar, M. R. *et al.* Text snippets to corroborate medical relations: An unsupervised approach using a knowledge graph and embeddings. In *AMIA Informatics Summit Proceedings*, vol. 2020 (American Medical Informatics Association, 2020).
27. Dastgheib, S. *et al.* Accelerating drug discovery in rare and complex diseases. In *International Semantic Web Conference (P&D/Industry/BlueSky)* (2018).
28. Proffitt, A. Pfizer's Model For The Intelligent Data Framework. <http://bit.ly/2JbShwv> (2019). [Online; accessed 19-July-2019].
29. Percha, B., Altman, R. B. & Wren, J. A global network of biomedical relationships derived from text. *Bioinformatics* **1**, 11 (2018).
30. Polleres, A., Kamdar, M. R., Fernandez Garcia, J. D., Tudorache, T. & Musen, M. A. A more decentralized vision for linked data. *Semantic Web* 1–19 (2019).
31. Wilkinson, M. D., Vandervalk, B. & McCarthy, L. The semantic automated discovery and integration (sadi) web service design-pattern, api and reference implementation. *Journal of biomedical semantics* **2**, 8 (2011).
32. Zaveri, A. & Ertaylan, G. Linked data for life sciences. *Algorithms* **10**, 126 (2017).
33. Debattista, J., Lange, C., Auer, S. & Cortis, D. Evaluating the quality of the LOD cloud: an empirical investigation. *Semantic Web* **9**, 859–901 (2018).
34. Kamdar, M. R., Zeginis, D., Hasnain, A., Decker, S. & Deus, H. F. ReVealD: A user-driven domain-specific interactive search platform for biomedical research. *Journal of biomedical informatics* **47**, 112–130, <https://doi.org/10.1016/j.jbi.2013.10.001> (2014).
35. Klyne, G. & Carroll, J. J. Resource description framework (RDF): Concepts and abstract syntax. *W3C recommendation* (2006).
36. McBride, B. The resource description framework (RDF) and its vocabulary description language RDFS. In *Handbook on ontologies*, 51–65, [https://doi.org/10.1007/978-3-540-24750-0\\_3](https://doi.org/10.1007/978-3-540-24750-0_3) (Springer Berlin Heidelberg, 2004).

37. Bechhofer, S. OWL: Web ontology language. In *Encyclopedia of Database Systems*, 2008–2009 (Springer, 2009).
38. Gruber, T. R. Toward principles for the design of ontologies used for knowledge sharing? *International journal of human-computer studies* **43**, 907–928 (1995).
39. Prud'Hommeaux, E., Seaborne, A. *et al.* SPARQL query language for RDF. *W3C recommendation* **15** (2008).
40. Kamdar, M. R. & Musen, M. A. PhLeGrA: Graph analytics in pharmacology over the web of life sciences linked open data. In *Proceedings of the 26th World Wide Web Conference, WWW 2017, Perth* (2017).
41. Marshall, M. S. *et al.* Emerging practices for mapping and linking life sciences data using RDF-A case series. *Web Semantics: Science, Services and Agents on the World Wide Web* **14**, 2–13 (2012).
42. Yates, A. *et al.* Ensembl 2016. *Nucleic acids research* **44**, D710–D716 (2016).
43. Maglott, D., Ostell, J., Pruitt, K. D. & Tatusova, T. Entrez gene: gene-centered information at ncbi. *Nucleic acids research* **33**, D54–D58 (2005).
44. Gray, K. A., Yates, B., Seal, R. L., Wright, M. W. & Bruford, E. A. Genenames.org: the HGNC resources in 2015. *Nucleic acids research* **43**, D1079–D1085 (2015).
45. Linked Open Vocabularies (LOV). <https://lov.linkeddata.es/dataset/lov/> (accessed October 09, 2019).
46. Whetzel, P. L. *et al.* BioPortal: enhanced functionality via new web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. *Nucleic acids research* **39**, W541–W545, <https://biportal.bioontology.org/> (2011).
47. Hastings, J. *et al.* The chemical information ontology: Provenance and disambiguation for chemical data on the biological semantic web. *PLOS ONE* **6**, 1–13, <https://doi.org/10.1371/journal.pone.0025513> (2011).
48. Kamdar, M. R., Tudorache, T. & Musen, M. A. A systematic analysis of term reuse and term overlap across biomedical ontologies. *Semantic web* **8**, 853–871 (2017).
49. Zaveri, A. & Dumontier, M. Ontology mapping for life science linked data. In *BMDID@ ISWC* (2016).
50. Haller, A., Fernández, J. D., Kamdar, M. R. & Polleres, A. What are links in linked open data? a characterization and evaluation of links between knowledge graphs on the web. *Working Papers on Information Systems, Information Business and Operations* (2019).
51. Hu, W., Qiu, H. & Dumontier, M. Link analysis of life science linked data. In *International Semantic Web Conference*, 446–462 (Springer, 2015).
52. Böhm, C. *et al.* Profiling linked open data with proLOD. In *Data Engineering Workshops (ICDEW), 2010 IEEE 26th International Conference on*, 175–178 (IEEE, 2010).
53. Hasnain, A. *et al.* A roadmap for navigating the life sciences linked open data cloud. In *Semantic Technology*, 97–112 (Springer, 2014).
54. Spahiu, B., Porrini, R., Palmonari, M., Rula, A. & Maurino, A. ABSTAT: ontology-driven linked data summaries with pattern minimalization. In *International Semantic Web Conference*, 381–395 (Springer, 2016).
55. Mihindukulasooriya, N., Poveda-Villalón, M., García-Castro, R. & Gómez-Pérez, A. Loupean online tool for inspecting datasets in the linked data cloud. In *International Semantic Web Conference (Posters & Demos)* (2015).
56. Hees, J., Bauer, R., Folz, J., Borth, D. & Dengel, A. An evolutionary algorithm to learn SPARQL queries for source-target-pairs. In *European Knowledge Acquisition Workshop*, 337–352 (2016).
57. Hewett, M. *et al.* PharmGKB: the pharmacogenetics knowledge base. *Nucleic acids research* **30**, 163–165 (2002).
58. Davis, A. P. *et al.* The comparative toxicogenomics database: update 2013. *Nucleic acids research* **41**, D1104–D1114 (2013).
59. Willighagen, E. L. *et al.* The ChEMBL database as linked open data. *Journal of cheminformatics* **5**, 23, <https://doi.org/10.1186/1758-2946-5-23> (2013).
60. Jupp, S. *et al.* The EBI RDF platform: linked open data for the life sciences. *Bioinformatics* **30**, 1338–1339 (2014).
61. Waagmeester, A. *et al.* Using the semantic web for rapid integration of WikiPathways with other biological online data resources. *PLoS computational biology* **12**, e1004989, <https://doi.org/10.1371/journal.pcbi.1004989> (2016).
62. Stearns, M. Q., Price, C., Spackman, K. A. & Wang, A. Y. SNOMED clinical terms: overview of the development process and project status. In *Proceedings of the AMIA Symposium*, 662 (American Medical Informatics Association, 2001).
63. Gil, Y. *et al.* PROV model primer. W3C Working Group Note (2013).
64. Isaac, A. & Summers, E. SKOS simple knowledge organization system primer. *Working Group Note*, W3C (2009).
65. Guha, R. V., Brickley, D. & Macbeth, S. Schema.org: evolution of structured data on the web. *Communications of the ACM* **59**, 44–51 (2016).
66. Kamdar, M. R., Tudorache, T. & Musen, M. A. Investigating term reuse and overlap in biomedical ontologies. In *Proceedings of the 6th International Conference on Biomedical Ontology, ICBO*, 27–30 (2015).
67. US National Library of Medicine. MEDLINE. <https://www.nlm.nih.gov/bsd/medline.html>. Accessed: 2019-06-09.
68. Pennington, J., Socher, R. & Manning, C. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543 (2014).
69. Kamdar, M. Biomedical word vectors. *figshare* <https://doi.org/10.6084/m9.figshare.9598760.v1> (2019).
70. Barrett, T. *et al.* BioProject and BioSample databases at NCBI: facilitating capture and organization of metadata. *Nucleic acids research* **40**, D57–D63 (2012).
71. Gonçalves, R. S., Kamdar, M. R. & Musen, M. A. Aligning biomedical metadata with ontologies using clustering and embeddings. In *European Semantic Web Conference*, 146–161 (Springer, 2019).
72. Blondel, V. D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* **2008**, P10008 (2008).
73. Clauset, A., Newman, M. E. & Moore, C. Finding community structure in very large networks. *Physical review E* **70**, 066111 (2004).
74. Markovitch, O. & Krasnogor, N. Predicting species emergence in simulated complex pre-biotic networks. *PLoS one* **13**, e0192871 (2018).
75. De Meo, P., Ferrara, E., Fiumara, G. & Provetti, A. Generalized Louvain method for community detection in large networks. In *Intelligent Systems Design and Applications (ISDA), 2011 11th International Conference on*, 88–93 (IEEE, 2011).
76. Shannon, P. *et al.* Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research* **13**, 2498–2504 (2003).
77. Biron, P. V., Malhotra, A., Consortium, W. W. W. *et al.* XML schema part 2: Datatypes (2004).
78. Gkoutos, G. V., Schofield, P. N. & Hoehndorf, R. The Units ontology: a tool for integrating units of measurement in science. *Database* **2012** (2012).
79. Kunze, J. & Baker, T. The Dublin core metadata element set. Tech. Rep., RFC 5013, August 2007).
80. Maali, F., Erickson, J. & Archer, P. Data catalog vocabulary (DCAT). *W3c recommendation* **16** (2014).
81. Skrbo, A., Begovic, B. & Skrbo, S. Classification of drugs using the ATC system (Anatomic, Therapeutic, Chemical Classification) and the latest changes. *Medicinski arhiv* **58**, 138–141 (2004).
82. Dumontier, M. *et al.* The SemanticScience Integrated Ontology (SIO) for biomedical research and knowledge discovery. *Journal of biomedical semantics* **5**, 14 (2014).
83. Simperl, E. Reusing ontologies on the Semantic Web: A feasibility study. *Data & Knowledge Engineering* **68**, 905–925, <https://doi.org/10.1016/j.datak.2009.02.002> (2009).
84. Corcho, O., Fernández-López, M. & Gómez-Pérez, A. Methodologies, tools and languages for building ontologies. Where is their meeting point? *Data & knowledge engineering* **46**, 41–64, [https://doi.org/10.1016/S0169-023X\(02\)00195-7](https://doi.org/10.1016/S0169-023X(02)00195-7) (2003).

85. Hamosh, A., Scott, A. F., Amberger, J. S., Bocchini, C. A. & McKusick, V. A. Online mendelian inheritance in man (OMIM), a knowledgebase of human genes and genetic disorders. *Nucleic acids research* **33**, D514–D517 (2005).
86. Bushman, B., Anderson, D. & Fu, G. Transforming the medical subject headings into linked data: creating the authorized version of MeSH in RDF. *Journal of library metadata* **15**, 157–176 (2015).
87. Wilkinson, M. D. *et al.* The fair guiding principles for scientific data management and stewardship. *Scientific data* **3**, 1–9 (2016).
88. Zaveri, A. *et al.* Quality assessment for linked data: A survey. *Semantic Web* **7**, 63–93 (2016).
89. Beek, W., Rietveld, L., Schlobach, S. & van Harmelen, F. Lod laundromat: Why the semantic web needs centralization (even if we don't like it). *IEEE Internet Computing* **20**, 78–81 (2016).
90. Vandenburg, P.-Y., Umbrich, J., Matteis, L., Hogan, A. & Buil-Aranda, C. SPARQLS: Monitoring public SPARQL endpoints. *Semantic Web* **8**, 1049–1065, <https://doi.org/10.3233/SW-170254> (2017).
91. Ontotext. Linked Life Data. <http://linkedlifedata.com/about>. Accessed: 2019-06-09.
92. Côté, R. G., Jones, P., Apweiler, R. & Hermjakob, H. The ontology lookup service, a lightweight cross-platform tool for controlled vocabulary queries. *BMC bioinformatics* **7**, 97 (2006).
93. Kamdar, M. Extracted schemas from the life sciences linked open data cloud. *figshare* <https://doi.org/10.6084/m9.figshare.12402425> (2020).
94. Déraspe, M. *et al.* Making linked data SPARQL with the InterMine biological data warehouse. In CEUR Workshop Proceedings, vol. 1795 (2016).
95. Kawashima, S., Katayama, T., Hatanaka, H., Kushida, T. & Takagi, T. NBDC RDF portal: a comprehensive repository for semantic data in life sciences. *Database* **2018** (2018).
96. Stark, C. *et al.* BioGRID: a general repository for interaction datasets. *Nucleic acids research* **34**, D535–D539, <https://doi.org/10.1093/nar/gkj109> (2006).
97. Kerrien, S. *et al.* The intAct molecular interaction database in 2012. *Nucleic acids research* **40**, D841–D846 (2011).
98. Cerami, E. G. *et al.* Pathway Commons, a web resource for biological pathway data. *Nucleic acids research* **39**, D685–D690 (2010).
99. Piñero, J. *et al.* DisGeNET: a comprehensive platform integrating information on human disease-associated genes and variants. *Nucleic acids research* **45**, D833–D839, <https://doi.org/10.1093/nar/gkw943> (2017).
100. Lane, L. *et al.* neXtProt: a knowledge platform for human proteins. *Nucleic acids research* **40**, D76–D83, <https://doi.org/10.1093/nar/gkr1179> (2011).
101. Boyce, R. D. *et al.* Dynamic enhancement of drug product labels to support drug safety, efficacy, and effectiveness. *Journal of biomedical semantics* **4**, 5 (2013).

## Acknowledgements

The authors would like to dedicate this paper to the memory of Dr. Amrapali Zaveri, a staunch advocate for data quality issues and the development of crowd-sourcing and human-in-the-loop methods to improve the quality of several data sources on the Web, who passed away in January 2020. The authors would also like to acknowledge Michel Dumontier, Javier Fernández, Rafael Gonçalves, Matthew Horridge, Marcos Martinez, Csongar Nyulas, Axel Polleres, and Tania Tudorache for valuable feedback on this research. M.K. and M.M. were supported in part by grants GM086587, GM103316, and HG004028 from the US National Institutes of Health.

## Author contributions

M.K. conducted the meta-analysis of data sources across the Life Sciences Linked Open Data cloud. All authors reviewed and approved the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to M.R.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021